Measurement Error and Misspecification in Demand-Based Diversion Ratios

Patrick Sun

August 3, 2022

OEA Working Paper 53

Office of Economics and Analytics
Federal Communications Commission
Washington, DC 20554

Abstract: When analyzing a horizontal merger it is often important to determine the extent to which the products of the merging firms are close substitutes. A commonly-used measure to assess the closeness of substitution is *the diversion ratio*: the fraction of demand leaving a product due to a price increase that goes to a specific rival product. One method of estimating diversion ratios is to first estimate a demand system and then calculate the implied diversion ratios. When estimating a demand system, two issues arise. First, using incorrect data values leads to measurement error. Second, using an incorrect model of demand leads to specification error. Through simulated datasets and using a known demand system, I examine how these errors can bias the diversion ratio estimates and a related, preliminary estimate of competitive harm, the Gross Upward Pricing Pressure Index (GUPPI). I find that even moderate measurement error results in biases comparable to the biases in share-based estimates of diversion that underestimate diversion between similar products. Further, I find that model specification error can result in substantial bias in the diversion and GUPPI estimates resulting in either overestimates or underestimates, depending on the specific nature of the specification error. My results suggest that: (1) measurement error is a serious concern when using demographic data to proxy for differences in price-sensitivity if that data do not very accurately represent the sample demographics; (2) practitioners should prefer flexible random coefficient models to avoid specification error; (3) practitioners should consider using observed markups for GUPPIs instead of estimating them; and (4) practitioners should more carefully consider the use of demand-based diversion ratios relative to alternatives..

Measurement Error and Misspecification in Demand-Based Diversion Ratios

1. Introduction

When analyzing a horizontal merger it is often important to determine the extent to which the products of the merging firms are close substitutes. A commonly-used measure to assess the closeness of substitution is *the diversion ratio* which is emphasized in the 2010 Horizontal Merger Guidelines.[1] The diversion ratio from an “origin” product to a “destination” product is defined as the share of quantity demanded of the origin product that is lost due to an increase in its price that is captured by the destination product. Diversion ratios, along with gross margin estimates, provide information on the likely incentive of the merged firm to raise prices post-merger, because a higher diversion ratio means that a higher percentage of lost sales caused by the price rise will go to the merging partner’s products. While estimates of diversion ratios based on demand system estimates are increasingly used in practical antitrust applications, little attention has been paid to the degree to which these estimates may be biased by measurement error and demand system misspecification. I examine the effect of measurement error and demand system misspecification using Monte Carlo experiments. I find that in certain situations both can lead to diversion ratio bias that is comparable to or worse than market share-based diversion ratios, , which are accurate for only very specific demand systems.

There are three main ways of estimating diversion ratios. First, diversion ratios can be directly calculated from data on consumer switching between products. One drawback of this approach is that, while firms often have data on gained or lost sales, they often do not have data on the specific products to which or from which consumers switch.[2] Further, such data often do not include why the consumers switched.[3] If consumers switch for reasons other than a price change, the implied diversion would not necessarily reflect what would happen after a price change, a major subject of interest in a merger review.[4]

Second, diversion ratios can be calculated as the market share of the destination product divided by one minus the share of the origin product. This “market share-proportional” or “market share-based” diversion ratio is easy to calculate from readily available market share data. It suffers from the problem that it is only a good estimate of diversion ratios if product switching happens in proportion to the product

---

[1] Horizontal Merger Guidelines, U.S. Department of Justice and the Federal Trade Commission at § 6.1 (Aug. 19, 2010). *See generally* Carl Shapiro & Howard Shelanski, *Judicial Response to the 2010 Horizontal Merger Guidelines*, 58 Rev. Indus. Org. 51 (2021) and Tommaso Valletti & Hans Zenger, *Mergers with Differentiated Products: Where Do We Stand?* 58 Rev. Indus. Org. 179 (2021).

[2] Consulting company Oxera notes that data for calculating diversion ratios directly “is rarely complete or available in an appropriate form.” Oxera, *Diversion Ratios: Why Does It Matter Where Customers Go if a Shop Is Closed?* Agenda: Advancing Economics in Business (Feb. 15. 2009), https://www.oxera.com/wp-content/uploads/2018/07/Diversion-ratios-updated_1-1.pdf-1.pdf. For issues regarding estimating diversion ratios from survey data, *see generally* Kirsten Edwards, *Estimating Diversion Ratios: Some Thoughts on Customer Survey Design*, *in* European. Competition Law Annual 2010: Merger Control in European and Global Perspective 31 (eds. Philip Lowe and Mel Marquis 2013).

[3] Knowing why consumers switched would require a survey in which a question was asked about why they switched, or “experimental” variation, where switching occurred after change in price when no other factors substantially changed.

[4] *See* Yongmin Chen & Marius Schwartz, *Churn Versus Diversion In Antitrust: An Illustrative Model*, 83 Economica 564 (2016).

choices of the market as a whole.[5] An example of such a case would be a demand system in which consumers choose products with certain probabilities, and these probabilities are the same for every single consumer. In contrast, if the consumers have differing tastes causing probabilities to vary across consumers, then market share-based diversion ratios likely would be a poor fit.

Third, diversion ratios can be estimated by first estimating a demand system and then mathematically deriving the implied diversion ratios. I refer to this as "demand-based diversion ratio estimation." Demand-based diversion ratios do not require switching data, and if demand is estimated well, the diversion ratios will capture substitution patterns that cannot be captured by market share-based diversion.[6] Some examples in merger review include the applicants' analysis in AT&T-DirecTV,[7] the U.S. government's analysis in Aetna-Humana,[8] and the applicants' analysis in T-Mobile-Sprint.[9]

Using Monte Carlo experiments, I document biases in demand-based diversion ratio estimates derived for a variety of commonly used demand systems that suffer from either measurement error or misspecification. Measurement error in the demand system naturally affects the estimated diversion ratios. Error in the specification will not only result in inaccurate estimation of demand, but also introduce errors in the formulas for implied diversion ratios. Although there is literature on the impact of specification error on demand estimation[10] and calibrated merger simulations,[11] I have been unable to find any paper that investigates how errors in the estimated demand system affect estimated diversion ratios.[12] Using a simulated set of demand data generated from a known demand system, I estimate diversion ratios using different demand models that may plausibly be used in a merger review. I compare the performance of these demand systems in diversion ratio estimation and in the estimation of the most

[5] Robert D. Willig, *Merger Analysis, Industrial Organization Theory, and Merger Guidelines*, 1991 Brookings Papers Econ. Activity: Microeconomics, 281, 301-04.

[6] Another potential benefit is that, if market definition is somewhat unclear, the demand system may be less sensitive to including too many products, a problem that can be significant in the case of market share-base diversion ratios. If the demand estimation is precise and accurate, products with little to no relation to the market in question will be estimated to have little substitution with the irrelevant products. However, this relies on overcoming all the difficulties faced in demand estimation, which may be compounded by adding irrelevant products. For example, Conlon & Mortimer (2013) find bias occurs when estimating demand on data in which inventories are commonly exhausted but all products are always assumed available. Christopher T. Conlon & Julie Mortimer, *Demand Estimation Under Incomplete Product Availability*, 5 Amer. Econ. J.: Microecon. 1 (2013).

[7] Letter from Maureen Jeffreys, Counsel to AT&T Inc., to Marlene H. Dortch, Secretary, FCC, MB Docket No. 14-90, Attach. (files Jul. 17, 2014), https://ecfsapi.fcc.gov/file/7521680277.pdf.

[8] Ari D. Gerstle, Helen C. Knudsen, June K. Lee, W. Robert Majure, & Dean V. Williamson, *Economics at the Antitrust Division 2016–2017: Healthcare, Nuclear Waste, and Agriculture*, 51 Rev. Indus. Org. 515, 522-23 (2017).

[9] Letter from Nancy Victory, Counsel to T-Mobile, to Marlene H. Dortch, Secretary, FCC, WT Docket No. 18-197, Attach. A (filed Nov. 6, 2018) (T-Mobile/Sprint Expert Economic Analysis), https://ecfsapi.fcc.gov/file/11060648404338/Nov.%206%20Public%20SuppResponse.pdf.

[10] Dongling Huang, Christian Rojas, & Frank Bass, *What Happens when Demand is Estimated with a Misspecified Model?* 56 J. Indus. Econ. 809 (2008).

[11] Philip Crooke, Luke Froeb, Steven Tschantz, & Gregory J. Werden, *Effects of Assumed Demand Form on Simulated Post-Merger Equilibria*, 15(3) Rev. Indus. Org. 205 (1999).

[12] One related paper is Rossi, Whitehouse & Moore (2019) which compares switching data-based diversion ratios estimates based on hospital referrals to demand-based diversion ratios estimates. However, the authors assume no systematic estimation error in their demand-based diversion ratios. Instead, they use the demand-based diversion ratios as a benchmark to evaluate the performance of the referral-based diversion ratios. Cecilia Rossi, Russell Whitehouse & Alex Moore, *Estimating Diversion Ratios In Hospital Mergers*, 15 J. Competition Law Econ., 639 (2019).

common measure of pricing pressure, the Gross Upward Pricing Pressure Index (GUPPI), which is the product of a diversion ratio and a ratio of the destination product markup over the origin product price.[13]

I find that even moderate levels of measurement error in data proxying for differences in price sensitivity can lead to significant biases. I use several specifications with different levels of measurement error, where the data used to measure the consumer's price sensitivity is only correlated with the true price sensitivity. Reducing the correlation to 0.75 produces diversion ratios that are similar to market share-based diversion ratios that use no consumer specific data at all and underestimates diversion between similar products. For comparison, consider the common practice of imputing an individual's data by using averages of their local area. If one imputes household income available in the 2019 American Community Survey (ACS) using the weighted median of the household's local area, the true income and the imputation have only a 0.32 correlation.[14] Further, a specification using data on five ordered groupings of consumer price sensitivities rather than the exact price sensitivity can lead to a numerically unstable estimate because some consumers' estimated demand has a positive price effect. In contrast, I find using ten groupings of price sensitivities yields more accurate and stable estimates, but this higher level of granularity is uncommon in practice.

The results examining misspecification suggest that incorrectly assuming the nature of price sensitivity leads to significant biases. For example, one specification is mis-specified by having no differences among consumers in their price sensitivities but captures differences in product value through a term that is specific to consumer type, product and market. This specification simultaneously predicts expected individual demand very well, while overestimating diversion originating from high price products because it overestimates the price sensitivity of consumers of high price products. A random coefficient specification, which takes account of variation in price sensitivity as a random variable, is accurate when it assumes price sensitivity has its true distribution, but has similar biases to the five ordered groupings specification when it uses a commonly assumed alternate distribution (i.e. assuming the distribution is normal when it is actually lognormal). The commonly used nested logit specification makes diversion reflect assigned product category shares. However, the estimated biases do not always improve upon, and are sometime worse than, the biases from share-based diversion.

These results suggest caution when using demand-based diversion ratio estimates. Unless the consumer demographic data accurately represent the true sample demographics, such as is typical with consumer panelist data or administrative data, measurement error will be a serious concern when using such data to proxy for difference in price sensitivities.. Whether you have consumer demographics data or not, misspecification is always a concern because matching the true variation in price sensitivity is crucial. The estimates with the higher and more imprecise estimates comes from inaccurate predictions that many consumers have positive price sensitivities, suggesting restriction of all price sensitivities to be

[13] Steven C. Salop & Serge Moresi, *Updating the Merger Guidelines: Comments*, 2009 Georgetown Law Faculty Publications and Other Works 1662, https://scholarship.law.georgetown.edu/cgi/viewcontent.cgi?article=2675&context=facpub.

[14] I use the 2019 One-Year ACS Public Use Microdata Sample (PUMS). PUMS is a subsample of ACS responses which meant to representative of national demographics and represents about 1% of the U.S, population. I use weighted medians of Public Use Microdata Areas (PUMAs), which are the smallest areas identified in the PUMS. PUMAs are designed to have approximately 100,000 residents. American Community Survey Office, United States Census Bureau, *AMERICAN COMMUNITY SURVEY 2019 ACS 1-YEAR PUMS FILES ReadMe* 3-4 (2020), https://www2.census.gov/programs-surveys/acs/tech_docs/pums/ACS2019_PUMS_README.pdf. For access to the data, use the "Explore Census Data" page. United States Census Bureau, *Explore Census Data*, https://data.census.gov/cedsci/ (last visited Nov. 19, 2021).

negative in the demand estimation. Practitioners could avoid these issues by using random coefficient specifications where price sensitivity distribution is flexibly estimated, though their applicability is limited due to increased complexity and data requirements. Measurement error and misspecification also impact markups estimated from a demand system, so practitioners should consider using observed markup data when available and trustworthy. In general, the challenges posed by errors in demand estimation suggest that practitioners should carefully consider the use of demand-based diversion ratios compared with alternative ways to estimate diversion ratios..

2. Diversion Ratios and Upward Pricing Pressure

Diversion as a tool for antitrust analysis appears as early as Willig (1991)[15] and the 1992 Horizontal Merger Guidelines,[16] with the DOJ using the term "diversion ratio" by 1995.[17] Shapiro (1996) presents an early treatment of diversion ratios as one of the U.S. antitrust authorities' tools.[18] Werden (1996) discusses how one can use diversion ratios to measure cost efficiencies necessary to offset merger-induced price-increases; he views this approach as a substitute for merger simulations based on parametric demand estimation, which he says are "vulnerable to attack" due to the need for functional form assumptions.[19] A later set of papers develops pricing pressure indices, which included diversion ratios as part of their formulas, and proposed these indices as an initial screen of likely anticompetitive effects in merger reviews.[20]

To explain the importance of diversion ratios, let us assume a differentiated products market t with $j \in J^t$ products. Firms charge all consumers the same price P_{jt} for product j in market t. Define the vector of all prices in market t as $\boldsymbol{P_t}$. I make no assumptions on the functional form for demand $D_{jt}(\boldsymbol{P_t})$ for each product, but for simplicity, I assume constant returns to scale costs for each product, where C_{jt} is the per unit cost for product j in market t.[21] $ftJ^{ft} \subset J^t f$

$$\pi_{ft} = \sum_{j \in J^{ft}} \left(P_{jt} - C_{jt}\right) D_{jt}(\boldsymbol{P_t}) \,. \tag{1}$$

The first order condition with respect to the price, P_{jt}, is

[15] Willig, *supra* note 5, at 299-305.
[16] Horizontal Merger Guidelines, U.S. Department of Justice and the Federal Trade Commission at § 2.2 (April 2, 1992).
[17] Carl Shapiro, *The 2010 Horizontal Merger Guidelines: From Hedgehog to Fox in Forty Years*, 77 Antitrust L.J. 701, 713-4 (2010).
[18] Carl Shapiro, *Mergers with Differentiated Products*, 10 Antitrust 23 (1996).
[19] Gregory J. Werden, *A Robust Test for Consumer Welfare Enhancing Mergers Among Sellers of Differentiated Products*, 44 J. Indus. Econ. 409 (1996).
[20] Daniel P. O'Brien & Steven C. Salop, *Competitive Effects of Partial Ownership: Financial Interest and Corporate Control*, 67 Antitrust L.J. 559 (1999); Salop & Moresi, *supra* note 13; and Joseph Farrell & Carl Shapiro, *Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition*, 10 BE J. Theoretical Econ. Article number: 000010220219351704156 3 (2010).
[21] The resulting equations are more complicated with non-constant marginal costs, where the shape of the supply curve becomes an important consideration in addition to diversion ratios. For most applications, researchers and practitioners assume constant marginal costs, which is a good approximation for industries at scale.

$$0 = D_{jt}(\boldsymbol{P_t}) + \left(P_{jt} - C_{jt}\right)\frac{\partial D_{jt}}{\partial P_{jt}} + \sum_{k \neq j \in J^{ft}} (P_{kt} - C_{kt})\frac{\partial D_{kt}}{\partial P_{jt}}. \tag{2}$$

I define the price-induced diversion ratio from product j to product k in market t as

$$DR_t^{jk} = -\frac{\partial D_{kt}}{\partial P_{jt}} / \frac{\partial D_{jt}}{\partial P_{jt}}. \tag{3}$$

Rearranging (2) we can find a formula of the markup for product j as a function of diversion ratios:

$$P_{jt} - C_{jt} = -D_{jt}(\boldsymbol{P_t}) / \frac{\partial D_{jt}}{\partial P_{jt}} + \sum_{k \neq j \in J^{ft}} (P_{kt} - C_{kt}) DR_t^{jk}. \tag{4}$$

The term $(P_{kt} - C_{kt})DR_t^{jk}$ represents how a firm that owns product *j* benefits from acquiring product *k*. A high diversion ratio from *j* to *k* and a high mark up on *k* implies the merged firm will recapture a large amount of lost profits from sales of *j* due to a price increase of *j* through additional sales of product *k*. Thus, the higher is the diversion ratios between an owned product and an acquired product the greater is the incentive (all else equal) to raise price post-merger.

In a horizontal merger between firm f and a rival firm h, the set of owned products of the merged firm is larger than that of the individual firms. Accordingly, if firm h sells a set of products $J^{ht} \subset J^t$ in market t as well, then additional terms of $(P_k - C_k)DR^{jk}$ for each product from firm h will be added to the post-merger formula for (4):

$$P_{jt} - C_{jt} = -D_{jt}(\boldsymbol{P_t}) / \frac{\partial D_{jt}}{\partial P_{jt}} + \sum_{k \neq j \in J^{ft} \cup J^{ht}} (P_{kt} - C_{kt}) DR_t^{jk}. \tag{5}$$

This means that two firms with high diversion ratios will have a higher incentive to increase markups after merging. As a result, the formula $(P_k - C_k)DR^{jk}$ calculated with pre-merger markup and diversion ratios has become an index of pricing pressure on product *j* due the addition of product k in merger review. "Upward Pricing Pressure," or UPP has been defined as:[22]

$$UPP_t^{jk} = (P_{kt} - C_{kt})DR_t^{jk}. \tag{6}$$

Note that equation (5) is a function of equilibrium prices, so the true additional post-merger markup cannot be directly calculated with pre-merger markups or diversion ratios. One can also think of UPP as a measure of the minimum required cost efficiencies necessary to maintain pre-merger prices,[23] but again, this is not exact because (5) uses pre-merger costs instead of post-merger costs.[24] Even so, in practice pre-merger markups and diversion ratios are used, with the understanding that this estimate of UPP is not excessively biased when the post-merger equilibrium prices are not significantly different from the pre-merger prices. Moreover, a large UPP estimated using pre-merger data implies a large change in equilibrium prices after the merger. In Monte Carlo experiments, Miller, Ryan, Miller and Sheu (2017)

[22] O'Brien & Salop, *supra* note 20 and Farrell & Shapiro, *supra* note 20.
[23] Farrell & Shapiro at 9-11, *supra* note 20.
[24] A more exact measure of the required efficiencies to keep prices equal is Compensating Marginal Cost Reductions (CMCR), which also uses diversion ratios. Werden, *supra* note 19.

show the UPP performs well with log-concave demand systems and performs comparably well to merger simulations with mis-specified models.[25]

In practice, the scale of UPP is dependent on how high prices are in the industry under study – a UPP of 10 is much less concerning if the average price of a good is $10,000 compared to when the average price is $10. It is therefore common to calculate the "Gross Upward Pricing Pressure Index" or GUPPI, which is the UPP divided by the price of the origin product:[26]

$$GUPPI_t^{jk} = \frac{(P_{kt} - C_{kt})}{P_{jt}} DR_t^{jk}. \tag{7}$$

There is no hard-and-fast rule on what value of GUPPI indicates a harmful merger, but Salop, Moresi, and Woodbury (2010) suggest less than 0.05 as presumptively *not* anticompetitive and greater than 0.1 as presumptively anticompetitive.[27] The literature has extended the GUPPI into new forms to study other post-merger incentives. For example, the "vGUPPI" is used to evaluate vertical mergers.[28]

3. Why Use Diversions Ratios if You Have Estimated a Demand System?

Hausman (2011) comments that it seems pointless to use an estimated demand system to estimate diversion ratios for UPP, because estimated demand systems can deliver post-merger price changes directly.[29] As I have noted earlier, UPPs and GUPPIs calculated with pre-merger data are only indicative of price changes rather than actual predictions of what those changes would be. In contrast, a fully estimated demand system can be used to directly predict the post-merger prices by changing supply-side assumptions, as in Nevo (2000).[30] Given assumptions about the nature of competition both before and after the merger, and the extent and nature of any merger efficiencies, an analyst can solve for the prices that would be optimal for all competing firms given the demand under the new market structure.

There are a few reasons why calculating diversion ratios with a demand system may be desirable, especially in policy and/or legal contexts. First, merger applicants have presented demand-based diversion ratios in merger reviews before and will likely continue to do so. Understanding likely biases in these measures is thus important regardless of alternative measures. Second, diversion ratios can be more easily grasped by non-economists and can be more indicative of close substitutes than cross-price elasticities.[31] This is especially helpful when communicating substitution patterns to regulatory or

[25] Nathan H. Miller, Marc Remer, Conor Ryan, & Gloria Sheu, *Upward Pricing Pressure as a Predictor of Merger Price Effects*, 52 Int'l J. Indus. Org. 216 (2017).
[26] Salop & Moresi, *supra* note 13.
[27] Steven C. Salop, Serge Moresi, & John R. Woodbury, *Scoring Unilateral Effects with the GUPPI: The Approach of the New Horizontal Merger Guidelines*, CRA Competition Memo, Charles River Associates (2010).
[28] Serge Moresi & Steven C. Salop, *vGUPPI: Scoring Unilateral Pricing Incentives in Vertical Mergers*, 79 Antitrust L.J. 185 (2013).
[29] "Of course, if an econometric demand model had already been estimated, there seems little reason not to perform a merger simulation rather than an upward pricing pressure calculation." Jerry Hausman, *2010 Merger Guidelines: Empirical Analysis*, 2011 Working Paper, 4 & n. 12, https://economics.mit.edu/files/6603.
[30] Aviv Nevo, *Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry*, 31 RAND J. Econ. 395 (2000).
[31] Conlon and Mortimer (2021) provide the following example. Assume there are three substitute products k, k', and j. Product k has a cross-price elasticity with j of 0.03 and a market share of 0.1. Product k' has a cross-price elasticity with j of 0.01 but a market share of 0.35. After a 1% price increase in j, the number of switchers to k'

judicial officials who review mergers, but otherwise do not regularly deal with economic issues. Such officials are also becoming more familiar with diversion ratios as their use in merger reviews becomes more common. Third, the nature of the supply-side may not be obvious due to industry complexity or non-public business practices, so any merger simulation would be making strong supply-side assumptions. If the nature of the demand side is more obvious, it may be more credible to present demand-based diversion ratios and simply point out whether merging firms have strong substitutes or not. Fourth, diversion ratios can also act as a transparent check on the performance of full merger simulations based on the same demand system. As shown above, the formulas for simulated price changes are very closely related to diversion ratios, so if the diversion ratios are wrong, demand simulations are also likely wrong. Finally, as mentioned earlier, Miller, Remer, Ryan, and Sheu (2017) have shown that UPP performs well compared to merger simulation under misspecification.[32] However, this finding is conditional on observing diversion ratios, so the contribution of the present paper is to characterize what happens when diversion ratios are not perfectly observed.

4. Discrete Choice Demand

The most popular framework amongst economists for studying differentiated product demand is the discrete choice random utility model (RUM).[33] Generically, consumer i in a market t of population N_t has a utility for product j with the bipartite form:

$$U_{ijt} = \delta_{ijt} + \epsilon_{ijt}. \tag{8}$$

ϵ_{ijt} is the product taste shock, a random variable that represents the component of utility each consumer has for product j that cannot be explained by observable data. δ_{ijt} is the component of utility that consumer i has for product j that varies as a function of product characteristics and possibly consumer characteristics. Because consumer level variation in δ_{ijt} is generally assumed to vary with product characteristics, individual differences in δ_{ijt} correspond to taste variation in product characteristics. I will follow most industrial organization applications and assume U_{ijt} is indirect utility, so δ_{ijt} is a function of price. Thus, an example of individual-varying sensitivity to a product characteristic is income specific price sensitivity: $\delta_{ijt} = \beta_j - \alpha_i P_{jt}$, where β_j represents a constant non-price utility each consumer has for product j and α_i measures individual-specific sensitivity to price P_{jt} that is a function of income.

Each consumer chooses only the product with the highest utility – hence the "discrete choice" moniker. Thus, expected individual demand (integrating over the product taste shocks ϵ_{ij}) is simply the individual choice probability of i choosing j, $E[D_{ijt}] = S_{ijt}$. Each consumer has a type $y_i \sim F$ (possibly observed or degenerate) that is independent of the product taste shocks, ϵ_{ij}. For a particular market t, all consumers face the same products so consumers' indices δ_{ijt} vary only by y_i. This implies that expected

(~0.35% of the market) would be larger than the number of switchers to k (~0.30% of the market), even though the cross-price elasticity to k would be larger. Thus given large differences in market share, cross-price elasticities can be misleading in terms of the closeness of substitution. Diversion ratios do not suffer this issue as they directly report the number of switchers. Christopher T. Conlon & Julie Holland Mortimer, *Empirical Properties of Diversion Ratios*, (2021) RAND J. Econ., 697 & n. 8.

[32] Miller, Remer, Ryan, & Sheu, *supra* note 25.

[33] *See generally* Simon P. Anderson, Andre De Palma, & Jacques-Francois Thisse, Discrete Choice Theory of Product Differentiation (1992); Kenneth E. Train, Discrete Choice Methods with Simulation, (2d ed. 2009); and Steven T. Berry & Philip Haile, *Foundations of Demand Estimation*, in Handbook of Industrial Organization, Vol. 4, 1 (eds. Kate Ho, Ali Hortaçsu, & Alessandro Lizzeri 2021).

market share, S_{jt} (integrating over the product taste shocks ϵ_{ij} and consumer types y_i), is simply the expected value of S_{ijt}:

$$S_{jt} = \frac{E[D_{jt}(\boldsymbol{P_t})]}{N} = \int E[D_{jt}(\boldsymbol{P_t})]\partial F(y_i) = \int S_{ijt}\partial F(y_i). \tag{9}$$

In principle, ϵ_{ij} could be drawn from any distribution. For example, if ϵ_{ij} is multivariate normal then choice probabilities follow the multinomial probit, which has been used in notable applications in industrial organization.[34] Still more popular is assuming that ϵ_{ij} is distributed i.i.d. according to the Gumbel distribution, also known as the Type I Extreme Value distribution.[35] In contrast to probit, which has no analytical form for the probabilities, the Gumbel distribution assumption implies the "logit" or "softmax" choice probability equations.[36] Assuming an outside option $j = 0$ with utility normalized to 0, the choice probabilities are[37]

$$S_{ijt} = \frac{\exp(\delta_{ijt})}{1 + \sum_{k \in J} \exp(\delta_{ikt})}. \tag{10}$$

The expected diversion ratios of discrete choice demand can be written as weighted averages.[38] Using (9) to rearrange (3) shows that the expected diversion ratio is a function of the derivatives of the expected market shares, which implies they are also functions of the derivatives of choice probabilities:

$$E\left[DR_t^{jk}\right] = -\frac{\dfrac{\partial E\left[D_{jt}(\boldsymbol{P_t})\right]}{\partial P_{jt}}}{\dfrac{\partial E\left[D_{jt}(\boldsymbol{P_t})\right]}{\partial P_{jt}}} = -\frac{\dfrac{\partial S_{kt}}{\partial P_{jt}}}{\dfrac{\partial S_{jt}}{\partial P_{jt}}} = -\frac{\int \dfrac{\partial S_{ikt}}{\partial P_{jt}}\partial F(y_i)}{\int \dfrac{\partial S_{ijt}}{\partial P_{jt}}\partial F(y_i)}. \tag{11}$$

Given a large number of consumers in a market, the expected and realized diversion ratios should be close in value. For the remainder of this study, I will treat the "the expected diversion ratio" as a highly accurate estimate of the realized "diversion ratio," and I refer to them interchangeably unless specifically noted.

Introducing the term $\frac{\partial S_{ijt}}{\partial P_{jt}} / \frac{\partial S_{ijt}}{\partial P_{jt}}$ inside the integral of the numerator of (11) allows the expected market-level diversion ratio to be rewritten as the weighted average of "individual diversion ratios," DR_i^{jk}:

$$E\left[DR_t^{jk}\right] = -\int \omega_{ijt} DR_{it}^{jk} \partial F(y_i). \tag{12}$$

[34] *E.g.* Austan Goolsbee & Amil Petrin, *The Consumer Gains from Direct Broadcast Satellites and the Competition With Cable TV*, 72 Econometrica 351 (2004).
[35] The Gumbel distribution is single-peaked, asymmetric and unbounded above and below. It has a cumulative density function of $F(\epsilon) = \exp(-\exp(\epsilon))$. Train, *supra* note 33, at 34.
[36] *Id.* at 36-37.
[37] The choice probabilities are the expectation of choice over the distribution of ϵ_{ij}, but not consumer type y_i. So each consumer has their own value for their δ_{ij}, and choice probabilities are based on δ_{ij}.
[38] The following discussion relies on the exposition by Conlon and Mortimer (2021). Conlon & Mortimer, *supra* note 31, at 698-706.

Individual diversion ratios, DR_{it}^{jk}, are equal to consumer i's ratio of the change of choice probability of product j over the change choice probability for product k due to a change of product $j's$ price:

$$DR_{it}^{jk} = \frac{\partial S_{ikt}}{\partial P_{jt}} / \frac{\partial S_{ijt}}{\partial P_{jt}}. \tag{13}$$

To avoid confusion between individual diversion ratios and market-level diversion ratios, I will refer to individual diversion ratios as "individual diversion ratios," while "diversion ratios" will refer to market-level diversion ratios.

Weights measure how strongly each consumer is expected to substitute away from j:

$$\omega_{ijt} = \frac{\frac{\partial S_{ijt}}{\partial P_{jt}}}{\int \frac{\partial S_{ijt}}{\partial P_{jt}} \partial F(y_i)}. \tag{14}$$

(12), (13), and (14) show there are two channels through which errors can bias diversion ratios – either individual diversion ratios are wrong or the "weight," measuring how much individual i is expected contributes to the total change in the market share of j, ω_{ijt}, is wrong.[39] Individual diversion ratios and weights are thus important diagnostic metrics for the rest of this paper.

Logit individual diversion ratios and weights can be derived substituting the formula for choice probabilities (10) into (13) and (14). The logit choice probabilities exhibit the Independence of Irrelevant Alternatives (IIA) property as defined by Luce (1995): as long as the same two options are available in otherwise different choice sets, the ratio between the choice probabilities for those two options is always equal.[40] IIA makes the individual diversion ratio a function of *only* the choice probabilities:

$$DR_{it}^{jk} = \frac{S_{ikt}}{1 - S_{ijt}}. \tag{15}$$

Weights are a function of the choice probabilities and the derivatives of δ_{ijt} with respect to price, i.e., the price sensitivity:

$$\omega_{ijt} = \frac{S_{ijt}\left(1 - S_{ijt}\right)\frac{\partial \delta_{ijt}}{\partial P_{jt}}}{\int S_{ijt}\left(1 - S_{ijt}\right)\frac{\partial \delta_{ijt}}{\partial P_{jt}} \, \partial F(y_i)}. \tag{16}$$

Thus market-level diversion of logit demand systems is not only a function of weights and individual diversion ratios, but also of choice probabilities and price sensitivities. Further, (16) shows

[39] An analogous derivation exists for the *realized* diversion rather than the *expected* diversion. In that case, diversion ratios are still a weighted average with the form of (12), but realized individual demand D_{ij} replaces choice probabilities S_{ij} in (13) and (14). The (13) and (14) analogues are not differentiable everywhere, because individual demand is binary for discrete choice demand. I therefore use the expected demand formulation in the current application of discrete choice; however, the realized demand formulation could be useful in in applications with continuous demand.
[40] R. Duncan Luce, Individual Choice Behavior: A Theoretical Analysis 9 (1959).

that accurate diversion ratio estimation requires accurate estimation of the *joint* distribution choice probabilities and the price sensitivities. If the marginal distributions of choice probabilities are correct, then individual diversion ratios can be accurately estimated. However, if the marginal distribution of price sensitivities or their joint distribution with choice probabilities are wrong, then the weights will not be estimated correctly. Measurement error in information about consumer types y_i thus can bias market-level diversion ratios by leading to biased choice probability estimates and biased estimates of price sensitivities. Specifications that do not have data on consumer types could work if they manage to somehow recreate the joint distribution of choice probabilities and price sensitivities, even if predictions for individuals are wrong. However, misspecification not only can result in biased choice probabilities but also may lead to direct misspecification of $\frac{\partial \delta_{ijt}}{\partial P_{jt}}$. This suggests that misspecification is a particular threat to unbiased estimation of weights.

One important case to consider is when δ_{ijt} has no variation across consumers, what I will call the “simple logit” model.[41] Choice probabilities and weights are then constant across consumers in the same market, and (15) simplifies to the market share-based diversion ratios:

$$E\left[DR_t^{jk}\right] = \frac{S_{kt}}{1 - S_{jt}}. \tag{17}$$

As seen from the derivation, share proportional diversion ratios mean that consumers have no taste variation for product characteristics. This seems implausible as basic economic theory implies consumers should vary at least in price sensitivity by income. Further, because diversion is perfectly proportional to market share, there is no way for two low share products to have high diversion with each other. This is inconsistent with the plausible situation where a niche of similar products might be unpopular with the general populace but the niche has a customer base that primarily substitutes within the niche.

For the remainder of this study, I will calculate market shares, diversion ratios and derived statistics like GUPPIs by using their expectations – this assumes the samples of the Monte Carlo experiments are large enough so that the expectation of the diversion ratio is a good approximation for the diversion ratio itself.

5. Monte Carlo Experiments Set Up

I perform Monte Carlo experiments in the following steps:

1. ***Simulate Consumers:*** I draw 500 consumer types y_i from the distribution F(y_i) and assign these to consumers in a market t. I do this 10 times for a total of 5,000 consumers distributed across 10 markets. 500 consumers is a reasonable number of consumers to observe per market in a market research survey[42] and allows me to perform the computations in a reasonable period of time given my limited computing resources. I experimented with using more markets to create more variation in prices but I achieve reasonably precise estimates for

[41] The literature often refers to this specification as simply the “logit” or “multinomial logit” model. This could be confusing in this paper since all the models in this paper use logit-type taste shock errors. I therefore use the term “simple logit” to improve clarity, which I borrow from Train, *supra* note 33, at 50.

[42] For example, Goolsbee and Petrin (2004) reports about 30,000 households over 317 markets in their data, which is about 100 consumers per market. Goolsbee & Petrin, *supra* note 34, at 356.

the true specification at 10 markets. I therefore keep the number of markets low for faster estimation.

2. ***Simulate Consumer Product Choices:*** Given the assortment of products J_t and the prices P_{jt} in market t, I calculate choice probabilities S_{jt} of selecting products j assuming a true demand system (explained later in this section). I then use these probabilities to divide up the unit interval into lengths equal to the choice probabilities. To simulate the choices, I then draw from the uniform distribution, and assign the chosen product depending on what region the draw is in.[43]
3. ***Estimate Demand System Specifications:*** With this synthetic dataset of consumers, products, markets, prices, consumers types y_i, and choices, I estimate all my 11 demand specifications using maximum likelihood estimation (or simulated maximum likelihood if the specification has random coefficients). The likelihood function is based on choice probabilities of observed choices.
4. ***Calculate Diversion Ratios and GUPPIs:*** I calculate the associated expected market-level diversion ratios and GUPPIs for each demand system based on the parameter estimates for Market 1. Aggregation to the market level requires an empirical analogue of integration over consumer types. In the case of observed types, I simply sum over the sample of consumers in each market. For specifications with random coefficients, I simulate 1,000 copies of each consumer with a different price sensitivity drawn using the estimated distribution. I calculate choice probabilities, individual diversion ratios and weights for each copy, and then average over all 1,000 copies of the synthetic data to calculate market-level diversion and GUPPIs.
5. ***Repeat:*** I then repeat steps 1-4 100 times to make 100 synthetic dataset and 100 sets of estimates for each of the 11 demand system specifications. I report bias, standard deviation of errors, and the root mean square error (RMSE). These results allow me to measure the expected sampling error in practical applications.

[43] For example, if there were only product 0 and product 1 in market t and $S_{i0t} = \frac{1}{3}$ and $S_{i1t} = \frac{2}{3}$, then I would divide the unit interval into region 0, $\left[0, \frac{1}{3}\right]$, and region 1, $(\frac{1}{3}, 1]$. If I draw from the uniform and the draw is less than $\frac{1}{3}$, I would assign product 0. If the draw is more than $\frac{1}{3}$ then I would assign product 1.

Table 1: Assortments, Prices and Market Shares

Market	Assortment	Low Tier Price (1 & 2)	High Tier Price (3 & 4)	Low Tier Market Share (1 & 2)	High Tier Market Share (3 & 4)
1	*1, 2, 3, 4*	1.34	2.17	45.7%	42.3%
2	*1, 2, 3, 4*	1.47	2.38	46.6%	39.2%
3	*1, 2, 3, 4*	1.67	2.71	47.3%	35.0%
4	*1, 2, 3, 4*	1.20	1.95	43.9%	46.2%
5	*1, 2, 3, 4*	1.00	1.63	41.2%	51.7%
6	*1, 3, 4*	1.34	2.17	35.2%	51.5%
7	*1, 2, 3*	1.34	2.17	55.6%	32.5%
8	*1, 2*	1.34	2.17	-	80.0%
9	*3, 4*	1.34	2.17	87.9%	-
10	*1, 3*	1.34	2.17	44.6%	42.4%

I will provide further detail in the following sections. This set up is markedly different from most of the related demand estimation literature in that it focuses on individual level choice data instead of market share data, and there are no product-market level unobservables.[44] The focus of the demand estimation literature is different from mine in that the literature typically addresses demand estimation under limited data and endogeneity so eschews other possible sources of error. Analogously, I assume a generous data environment and no endogeneity to fully isolate the impact of measurement error and misspecification. This present paper should be seen as complementary to the pre-existing literature for informing practitioners about demand estimation.

5.1 Markets, Assortments, Qualities, Prices

I define 10 different markets, which will vary by the assortment of products from which consumers can choose and by the price levels of the products. Table 1 presents variation in products across markets. Market 1 will be my baseline market – it will include all possible products and will have my baseline prices. Markets 2-10 will be altered from Market 1 to aid in identifying parameters of the various demand systems. If I only used Market 1, it would not be possible to distinguish variation in utility due to price differences from variation due to quality differences. My specifications allow differences in both quality and prices across products but not within markets, so I need multiple markets with the same product to measure how price changes demand given the same quality.

Markets 1 to 5 will have 4 different products plus an Outside Option. There are two product quality tiers: symmetric products 1 and 2 of the "Low Tier" have quality $Q_j = 2.5$, while symmetric products 3 and 4 of the "High Tier" have $Q_j = 3$. Other markets vary by assortment to help with identification of the model parameters that govern substitution across products: comparing markets with a

[44] *See generally* Berry & Haile, *supra* note 33.

product and markets without that product shows how consumers will substitute when forced to switch.[45] Market 6 will not have Product 2; Market 7 will not have Product 4; Market 8 will not have High Tier products; Market 9 will not have Low Tier Products; and Market 10 will not have Products 2 or 4.

For simplicity, I assume firms only produce one product. Firms produce with constant marginal cost. Firm *f*'s profits from producing its one product j are equal to the one-product version of equation (1):

$$\pi_{ft} = \left(P_{jt} - C_{jt}\right)D_{jt}. \tag{18}$$

In line with the traditional diversion ratios assumption of firm conduct and market structure, firms compete as Nash-Bertrand price setters in Market 1. Market 1 prices are thus the equilibrium strategies of the Nash-Bertrand price setting game of the firms selling to the entire Market 1 population of all 100 datasets. This is analogous to treating the full set of simulations as the "population" of the market and each individual dataset as a sample.[46] High Tier products cost more to produce than Low Tier Products: in market 1, Low Tier products have marginal costs of $C_j = 1.0$ and High Tier products have marginal costs of $C_j = 1.5$.

I directly assume the variation in the prices for Markets 2 to 5 to aid in identifying price sensitivity parameters. This price variation can be thought as the result of cost variation across markets or some regulatory intervention. I eschew modeling how exactly the supply side works in these other markets because finding the equilibrium costs that would imply these exact prices can be complex. Relative to prices in Market 1, Market 2 prices are 10% greater, Market 3 prices are 25% greater, Market 4 prices are 10% lower, and Market 5 prices are 25% lower. For the same reason, I keep the prices in Markets 6-10 the same as Market 1 as the demand under different assortments can be directly compared to demand in Market 1 without potentially confounding variation in prices.. The resulting prices can be observed in Table 1.

Using all markets to calculate diversion and GUPPIs introduces composition effects in aggregate diversion ratios and GUPPI estimates.[47] I therefore use only Market 1 data when deriving diversion ratios and GUPPIs for simpler interpretation of the results. I use data for Markets 2-10 only in demand estimation.

5.2 The "True" Demand System

I will assume that consumers in our simulations have "true" discrete choice demand with logit product taste shocks and that δ_{ijt} has the form

[45] However, as the results will show, measurement error and mis-specification may prevent a specification from estimating why a consumer switched (i.e. based on their price sensitivity), and thus generate incorrect predictions for diversion due to a price change.
[46] There is an argument that it would have been more appropriate to construct our individual datasets via bootstrapping from the entire population of the 100 synthetic datasets, but this would have only given a negligible conceptual benefit at a significant increase in complexity.
[47] For example, if I include both Market 1 and Market 6 in a single diversion ratio calculation, I would need to account for the fact that half of consumers in Market 6 cannot divert to or from the missing Product 2. Inclusion of a market with different prices like Market 2 is feasible, but then the resulting diversion ratios would be a blend of diversion from two separate price levels.

$$\delta_{ijt} = \theta\big(Q_j - y_i P_{jt}\big). \tag{19}$$

Q_j is quality; P_{jt} is market t-specific price; and y_i is an individual-specific consumer type that (up to the factor θ) explains price sensitivity. In many demand models, a common example of such a characteristic would be income, which reduces price sensitivity through the income effect. In (19), the price sensitivity *increases* in y_i, so y_i is more analogous to inverse or negative income than income itself.

Conditional on the Outside Option's $\delta_{i0t} = 0$ as a normalization, larger θ implies that taste shock ϵ_{ij} is a smaller component of utility than δ_{ijt}. That is, if P_{jt} is measured in dollars, then the standard deviation of ϵ_{ij} is $\frac{\pi}{\sqrt{6}\,\theta y_i}$ dollars.[48] As θ approaches to 0, the model approaches a simple logit where $\delta_{ijt} = 0$ for all consumers, products and markets: products and the Outside Option are chosen at random with equal probability. As θ tends to infinity, the model approaches the "vertical model" of Shaked and Sutton (1982) and Bresnahan (1987), where all taste variation is explained by variation in price sensitivity and there are no product taste shocks.[49]

This demand system is intentionally very simple. While there are interesting questions about how the interaction between price sensitivity and the taste for other product characteristics impacts diversion ratios, using a simpler model will make the results more transparent. I therefore use only one random coefficient and also absorb all non-price characteristics into the single quality variable. Further, I assume that there is no unobserved product-market utility, i.e. price endogeneity is not an issue.

As the true model is still logit-based, (15) remains the formula for individual diversion ratios. With linear δ_{ijt}, price sensitivity is $\frac{\partial \delta_{ijt}}{\partial P_{jt}} = \theta y_i$, so there is a more specific version of the (16) formula for weights:

$$\omega_{ijt} = \frac{\big(1 - S_{ijt}\big) S_{ijt} y_i}{\int \big(1 - S_{ijt}\big) S_{ijt} y_i \partial F(y_i)}. \tag{20}$$

The formula for individual diversion ratios are typical for logit-based models, but y_i directly impacts the value of the weights. What is less obvious is the role of θ; high θ mechanically makes all δ_{ijt} larger in magnitude, so choice probabilities (relative to the choice probability for the Outside Option) are more extreme for extreme values of y_i. As choice probabilities appear in the weights, this can create a more lopsided weighted average.

I will assume y_i is lognormal distributed with a mean of 1 and skewness of 5.[50] I use the lognormal distribution because this guarantees that price sensitivity will always be negative. I choose a mean of 1 for ease of exposition: utility of 1 unit of price P_{jt} (i.e., money) to the mean consumer is θ. I choose a high skewness because this accentuates the differentiation of consumers across their chosen

[48] Relative to the units of the coefficients of an estimated multinomial logit, the standard deviation of taste shock ϵ_{ij} is $\frac{\pi}{\sqrt{6}}$. Train, *supra* note 33, at 40-41.

[49] Avner Shaked & John Sutton, *Relaxing Price Competition Through Product Differentiation*, 49 Rev. of Econ. Stud. 3, (1982); and Timothy F. Bresnahan, *Competition and Collusion in the American Automobile Industry: The 1955 Price War*, 35 J. Indus. Econ. 457 (1987).

[50] A lognormal random variable is the log of a normal random variable with mean μ and standard deviation σ. For this paper, the parameters underlying the normal distribution seem less meaningful than the mean and skewness of the lognormal distribution, but for interested readers, I am effectively assuming $\mu = -0.42$ and $\sigma = 0.92$.

products. Diversion ratios will then be very different from the case in which all consumers have the same price sensitivity, i.e., share-proportional diversion.

5.3 θ and GUPPIs

To limit the amount of information reported, I will report only one diversion within the Low Tier (Product 1 to Product 2), one diversion from the Low Tier to the High Tier (Product 1 to Product 3), one diversion from the High Tier to Low Tier (Product 3 to Product 1), and one diversion within the High Tier (Product 3 to Product 4). As products within tiers are symmetric in quality and price, the presented diversions are virtually identical (up to the simulation/estimation error) to the diversions I do not present. For example, the diversion ratio of Product 1 to Product 3 is nearly identical to the diversion ratio of Product 2 to Product 4 because they are both Low Tier to High Tier diversions.

Due to computational time, I focus on a single value of θ. However, it is instructive to show how varying θ impacts the diversion and GUPPIs, as this informs the θ that I ultimately select. For the Market 1 subset of all 100 datasets, I calculate the overall diversion ratios and GUPPIs for all consumers in all datasets. Figure 1 indicates that, as θ increases (i.e., making price sensitivity variation more important relative to product taste shocks), diversion within tiers goes up noticeably. However, Figure 2 indicates that the GUPPI decreases as θ increases. Recall that GUPPIs are the product of percent markups and diversion. As θ increases, own-price elasticity increases and thus markups decrease as seen in Figure 3. Markups decrease more rapidly than diversion increases, so the markup decrease dominates in their resulting product.

For my purposes I choose $\theta = 5$ as this implies a sizeable gap between true GUPPIs and share-proportional GUPPIs for both within-tier diversions. Table 1 shows market shares and prices for Market 1. Low Tier products have market shares of 23% and High Tier products have market shares of 21%. Low Tier products have prices of 1.37 (markup is 27% of price) and High Tier products have prices of 2.17 (markup is 31% of price). Table 2 shows Market 1 diversion ratios: expected diversion is high within tier: 59% in the Low Tier and 38% in the High Tier. This compares to the share-proportional diversion of 29% within the Low Tier and 27% in the High Tier. Diversion between the High Tier and the Outside Option is essentially 0, which means the model with its high skewness of price sensitivity is close to the “vertical model” where products only substitute to the next-lowest and next-highest products in quality level. For comparison, Table 3 shows share-proportional diversion: all products beside the Outside Option have roughly the same market share so diversion is close to symmetric.

Individual diversion ratios and weights differ greatly by products and y_i. Figure 4 shows that individual diversion is mostly dominated by the destination product of diversion. Diversion ratios to High Tier products are strongest amongst weakly price sensitive consumers with low y_i. Diversion ratios to the Outside Option are strongest amongst strongly price sensitive consumers with high y_i Diversion ratios to Low Tier products are strongest amongst moderate price sensitive consumers with moderate y_i. Figure 5 shows weights are also determined by price sensitivity. An extreme amount of weight is put on highly price sensitive consumers with large y_i for diversion from the Outside Option. Weight is highest for low y_i for the High Tier products, and the weight is highest for medium y_i for the Low Tier goods. In summary, consumers with high individual diversion ratios having high weights is what makes within-tier diversion high in my baseline specification.

Figure 1: Diversion Ratios in Market 1 as θ Changes

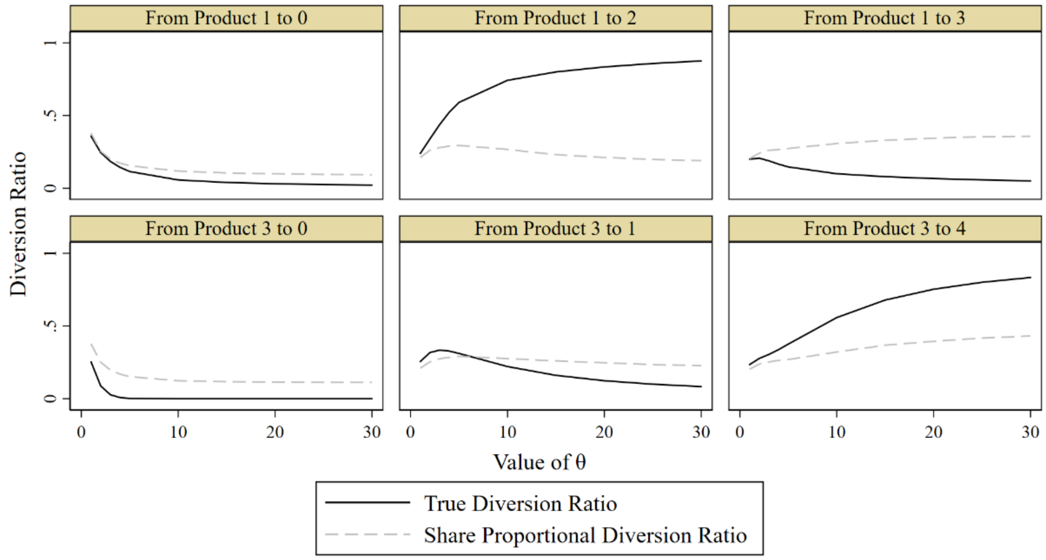

Evaluated at $\theta = \{1, 2, 3, 4, 5, 10, 15, 20, 25, 30\}$.

Figure 2: GUPPIs in Market 1 as θ Changes

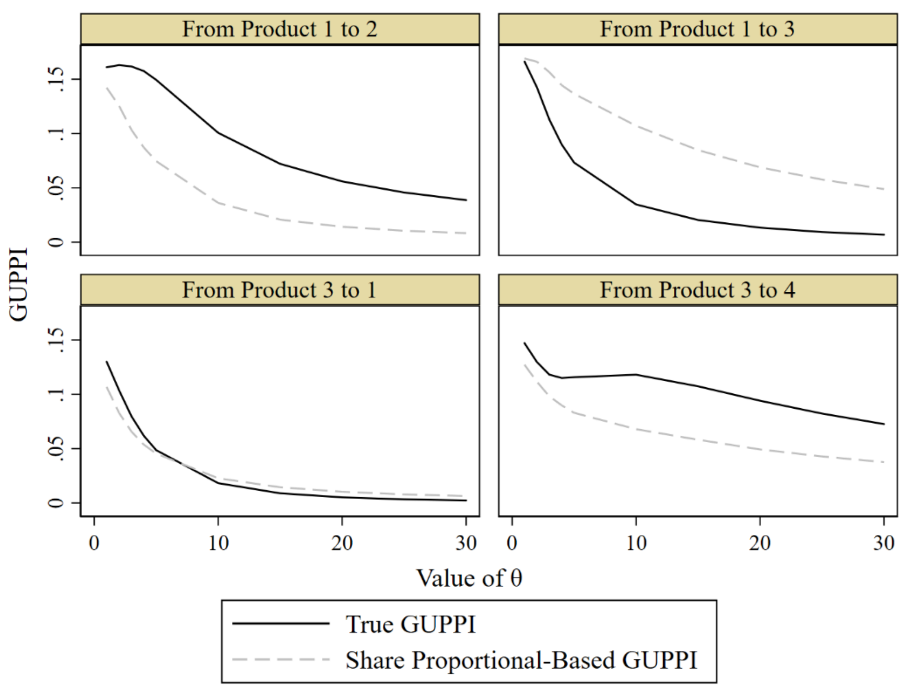

Evaluated at $\theta = \{1, 2, 3, 4, 5, 10, 15, 20, 25, 30\}$.

Figure 3: Markups in Market 1 as θ Changes

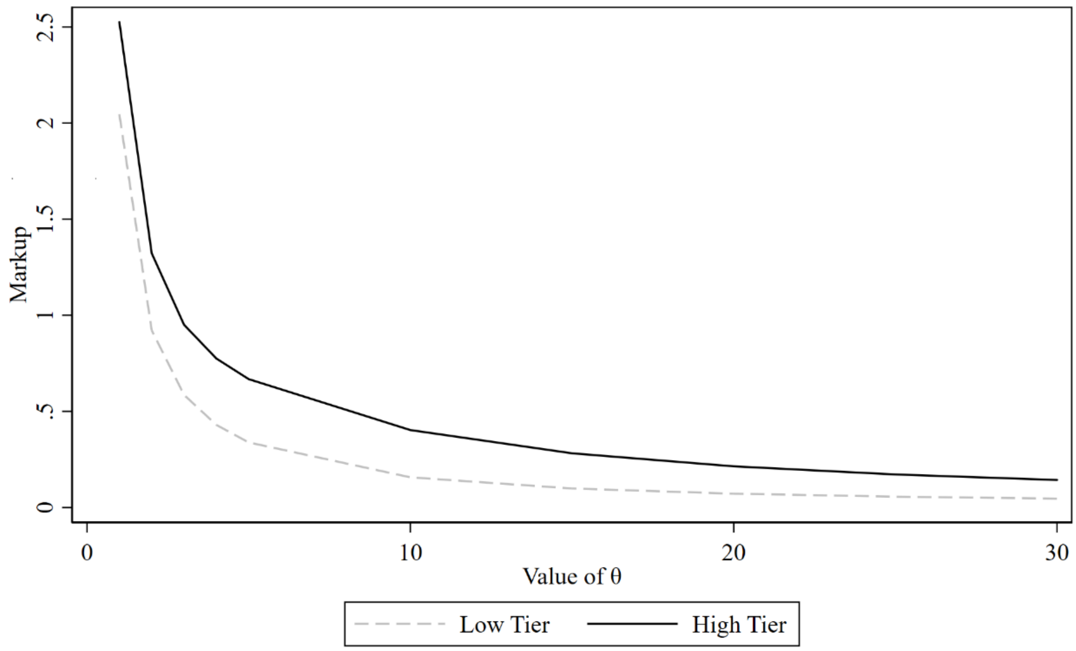

Evaluated at $\theta \in \{1, 2, 3, 4, 5, 10, 15, 20, 25, 30\}$.

Table 2: True Diversion Ratios (Market 1, $\boldsymbol{\theta = 5}$)

		Diverting to…				
		0	*1*	*2*	*3*	*4*
Diverting from…	*0*	-	0.496	0.496	0.004	0.004
	1	0.115	-	0.591	0.147	0.147
	2	0.115	0.591	-	0.147	0.147
	3	0.002	0.311	0.311	-	0.376
	4	0.002	0.311	0.311	0.376	-

Diversion ratios calculated over all simulated consumers in Market 1 across all 100 synthetic datasets .

Table 3: Estimated Share-Proportional Diversion Ratios (Market 1, $\boldsymbol{\theta = 5}$)

		Diverting to…				
		0	*1*	*2*	*3*	*4*
Diverting from…	*0*	-	0.261	0.258	0.240	0.241
	1	0.155	-	0.295	0.274	0.276
	2	0.155	0.297	-	0.273	0.275
	3	0.152	0.291	0.288	-	0.269
	4	0.152	0.291	0.289	0.268	-

Diversion ratios calculated for all consumers in Market 1 across all 100 synthetic datasets.

Figure 4: Individual Diversion Ratios by y_i Percentile (Market 1, $\theta = 5$)

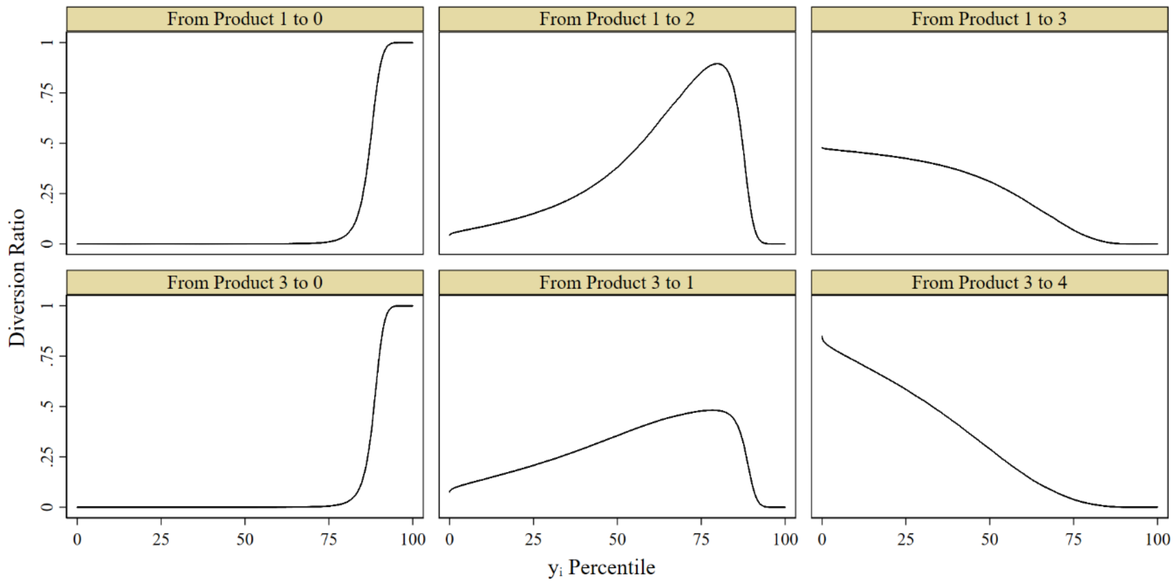

Individual diversion ratios calculated for all simulated consumers across all synthetic datasets in Market 1.

Figure 5: Individual Weights by y_i Percentile ($\theta = 5$)

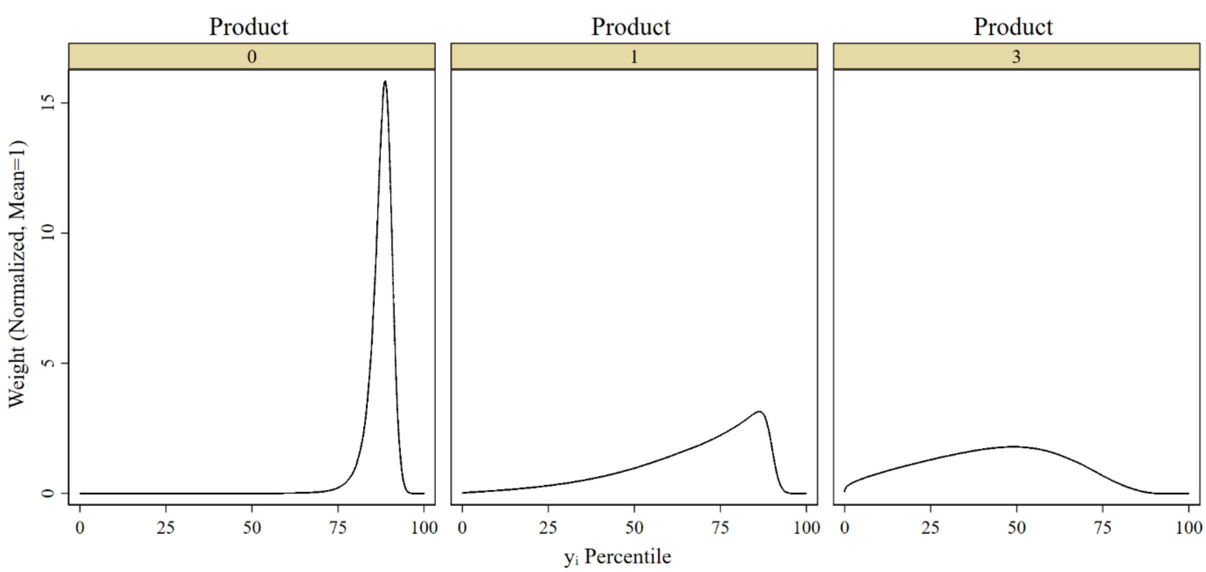

Weights calculated for all simulated consumers in Market 1 across all synthetic datasets. Weights are normalized so that the mean weight is 1.

Table 4: Prices Changes and GUPPIs (Market 1, $\theta = 5$, *Merging Product Bolded*)

Merger Type	Product	Pre-Merger Price	GUPPI	Post-Merger Price	Fraction Price Change
Low & Low	***1***	**1.34**	**0.15**	**1.50**	**0.12**
	2	**1.34**	**0.15**	**1.50**	**0.12**
	3	2.17	-	2.09	-0.04
	4	2.17	-	2.09	-0.04
Low & High	***1***	**1.34**	**0.07**	**1.42**	**0.06**
	2	1.34	-	1.37	0.02
	3	**2.17**	**0.05**	**2.29**	**0.06**
	4	2.17	-	2.17	0.00
High & High	*1*	1.34	-	1.41	0.06
	2	1.34	-	1.41	0.06
	3	**2.17**	**0.12**	**2.72**	**0.26**
	4	**2.17**	**0.12**	**2.72**	**0.26**

An important feature of the model is that the expected weights for products drops off quickly as y_i approaches zero because weights are proportional to y_i: less price sensitive consumers have lower weights than price sensitive consumers because they respond less to price. The High Tier consumers have disproportionally low y_i so marginal increases in the price of High Tier products result in more muted increases in unit sales lost than when prices increase for Low Tier Products. Given the market shares are all about equal for non-outside option products, market share-based diversion ratios between these products are roughly similar (27%-30%). Since the baseline of share-based diversion ratios is similar, when I compare the difference between true diversion and share-based diversion, the differences are more pronounced for the Low Tier goods that have more price sensitive consumers: the within-High Tier diversion ratio is substantially greater than the share-based diversion ratio (38% vs. 27%) but not as high as the difference for within-Low Tier diversion (59% vs. 30%). This also means specifications that do not properly account for y_i heterogeneity will have a tendency to overestimate diversion originating from higher tier products because they will overweight the individual diversion ratios of very low y_i.

Given these parameter assumptions, Market 1 is problematic for any merger. Table 4 shows GUPPIs and post-merger price changes based on the true demand system. All possible mergers will result in higher than 5% price increases for products of merging firms and, in the case of a merger of the High Tier products, all products.[51] Using the 0.05 and 0.10 thresholds for GUPPIs explained in the earlier section on GUPPIs[52], any merger within tier should be presumptively anticompetitive and a cross-tier merger should warrant further review: GUPPIs between the Low Tier products are 0.15 while GUPPIs between the High Tier products are 0.12; and Low-to-High GUPPIs are 0.7 while High-to-Low GUPPIs are 0.5. GUPPIs are similar in magnitude to the resulting price increases in Low-Low and Low-High

[51] Interestingly, a merger between Low Tier products generates slightly negative pricing pressure: the Low Tier price increases are so extreme that the residual demand for High Tier products becomes much more price sensitive so that High Tier prices end up slightly lower despite the reduction in competition.
[52] Salop, Moresi, & Woodbury, *supra* note 27.

mergers, but underestimate the 26% price increase in a High-High merger – the low price sensitivity of High Tier consumers support higher price hikes and make the post-merger equilibrium different enough that the GUPPI using pre-merger data are no longer accurate. This failure to perfectly replicate these price changes reinforces that GUPPIs are not direct estimates of price changes themselves.[53]

GUPPI calculations for the estimated demand specifications raise the issue of markup estimation. Markups are often themselves estimated in practice because marginal costs are not generally observed. One common estimation method for markups is to use accounting data, but accounting standards often conflate fixed and variable costs, and only deal with measurable pecuniary costs. Further, accounting cost data are often only available at the firm level, causing averaging errors for multiproduct firms.[54]

In academic research with full demand estimation, margins are often estimated using the first order condition of the profit function. In my setup of Nash-Bertrand competition between single product firms with discrete choice demand, expected markups are a version of equation (4) without cannibalization effects:

$$E\left[P_{jt} - C_{jt}\right] = -\frac{\int S_{ijt} \partial F(y_i)}{\int \frac{\partial S_{ijt}}{\partial P_{jt}} \partial F(y_i)}. \tag{21}$$

While expected demand $\int S_{ijt} \partial F(y_i)$ is likely to be estimated well or realized demand may be available, $\int \frac{\partial S_{ijt}}{\partial P_{jt}} \partial F(y_i)$, the denominator of (21) is also the denominator of the weight function in (14). Thus problems in estimating the weights will also affect estimation of markups. When I report GUPPIs, I will report both estimates using the true markups as well as those using estimated markups from the demand system.

6. Demand Specifications

For each dataset, I will estimate demand using specifications with precedent in antitrust analysis for comparison. I simulate product choices at the individual level, so I maximize the negative log-likelihood of the choices made by each simulated consumer. I report the mean and standard deviation of the estimated coefficients over the 100 datasets to give a sense of the precision of the estimates.

Below are the details of the specifications I estimate. They vary either by how accurately data on consumer type y_i is observed, which reflects measurement error, or by the functional form of δ_{ijt}, which reflects model misspecification.

---

[53] There is a lack of consensus on what UPPs and GUPPIs exactly represent. "UPP does not predict post-merger prices, but only predicts the sign of changes in price." Joseph Farrell & Carl Shapiro, *Upward Pricing Pressure In Horizontal Merger Analysis: Reply to Epstein and Rubinfeld*, 10 BE J. Theoretical Econ., 3 (2010). "We do see UPP as a simple and useful measure that is generally indicative of likely price effects." Joseph Farrell & Carl Shapiro, *Upward Pricing Pressure and Critical Loss Analysis: Response*, 2010 CPI Antitrust J., 4, https://www.competitionpolicyinternational.com/assets/Uploads/Shapiro-FarrellFEB10-copy.pdf

[54] *See generally* Seth B. Sacher & John Simpson. *Estimating Incremental Margins for Diversion Analysis*, 83 Antitrust L.J. 527 (2020).

6.1 Simple Logit Specification

My benchmark for poor performance is a simple logit specification estimated on choice and price data but no other individual-specific data. Recall that this is essentially assumed when an analyst uses share-proportional diversion ratios. I assume the δ_{ijt} for this specification is

$$\delta_{ijt}^{PURE} = \beta_j^{\mathrm{PURE}} + \alpha^{PURE} P_{jt}. \tag{22}$$

Given the lack of individual heterogeneity, it is not possible for this specification to predict distribution of the choice probabilities or the price sensitivities. Thus it is practically impossible for the specification to generate the correct diversion ratios because the weighted average of individual diversion ratios (12) would have to coincidentally equal the share proportional formula (15).

6.2 Correct Model and Data Specification

For this specification, I use the "correct" model of utility of equation (19) used to generate the synthetic data, and assume y_i is observed. This corresponds to the case where the econometrician has access to individual-specific data on consumer characteristics (micro-data) which is informative for product choice. In this case, the micro-data on y_i is perfectly informative of price sensitivity as price sensitivity is proportional to y_i. I use a common transformation to include the impact of varying price sensitivity: I include both price and interactions of price with a centered y_i and denote the centered y_i as $\dot{y}_i$. I use the following form of the indices for estimation:

$$\delta_{ijt}^{TRUE} = \beta_j^{\mathrm{TRUE}} + \alpha^{TRUE} P_{jt} + \gamma^{TRUE} \dot{y}_i P_{jt}. \tag{23}$$

It follows from (19) and the assumptions that $\theta = 5$ and the mean of y_i is 1 that this specification estimates the true choices probabilities and price sensitivities if estimated parameters $\hat{\beta}_j^{TRUE} = 5Q_j$, $\hat{\alpha}^{TRUE} = -5$, $\hat{\gamma}^{TRUE} = -5$. This specification allows decomposition of the price effect into its mean plus a varying term, which is common when the demand researcher is unsure whether price sensitivity varies.

There are drawbacks to this specification. The first is that the additional parameter makes estimation less efficient. The second is that for very negative $\dot{y}_i$ (near zero y_i), this can imply positive price sensitivity if $\alpha^{TRUE} > \gamma^{TRUE}$. The denominator of equation (20), the formula for weights, is a function of choice probabilities and price sensitivities. When price sensitivities are both positive and negative, the denominator of weights can be small and/or be the opposite sign of its numerator. Thus, some estimated weights can then be unrealistically large in magnitude and/or strongly negative. Moreover, it could be profit maximizing to charge extremely high prices. While most consumers may no longer buy such highly priced products, consumers with positive price sensitivities would continue to buy no matter what the price is. It therefore may be a corner solution to charge an incredibly high price to just these price-loving consumers rather than to sell to consumers more generally.[55] The diversion ratios would then no longer be informative of the post-merger equilibrium, because pricing equation (5) is an interior solution. In general, applications tend to ignore this issue, so I will simply assume that budget

[55] Under the assumption that U_{ijt} represents indirect utility, there are budget constraints so prices cannot be infinite.

constraints are such that no local corner maximum yields greater profit than the interior solution.[56] I will document these positive price coefficients when they occur in the estimation result for this and other specifications.

6.3 Correct Model but Mismeasured Data Specification

For three specifications, I use mismeasured y_i data to explore measurement error. Mismeasured data have the same marginal distribution of y_i, but are only correlated with y_i.[57] I use correlations of 0.95, 0.85, and 0.75 to illustrate the decline in performance as the data become less accurate. These specifications not only correspond to when measured y_i has been contaminated with errors, but also the common case when y_i is unobserved but some correlated proxy is used instead.

I denote the centered version of mismeasured y_i with correlation ρ with y_i as $\dot{y}_i(\rho)$. The associated choice probabilities are determined by the δ_{ijt}:

$$\delta_{ijt}^{MM\rho} = \beta_j^{\mathrm{MM\rho}} + \alpha^{MM\rho} P_{jt} + \gamma^{MM\rho} P_{jt} \dot{y}_i(\rho). \tag{24}$$

The resulting bias will depend on the degree of inaccuracy, but it is an open question on how severe the issue is in practice.

6.4 Quantile Coefficient Specification

In contrast to the case where mismeasured data have the correct marginal distribution, another common form of measurement error is where y_i is only observable up to discrete ranges, which I will call "bins" indexed by $b \in B$. For example, such data are often produced by surveys where giving a range is easier than giving an exact number. Binning is also used as a strategy of anonymizing data when releasing to the public.

To recreate this case, I estimate a model with bins designated by quantiles. The δ_{ijt}takes the form

$$\delta_{ij}^{QUANTILE} = \beta_j + \alpha_{QUANTILE}(y_i) P_{jt}. \tag{25}$$

$\alpha_{QUANTILES}(y_i)$ is a bin-specific coefficient which take a particular value based on what quantile y_i is in. Weights will take on the form

[56] For example, Nevo (2000) estimates a distribution with 0.7% positive price sensitivities. Aviv Nevo, *Measuring Market Power in the Ready-to-Eat Cereal Industry*, 69 Econometrica 307, 329 (2000).

[57] I generate a Z with correlation ρ with Y and the same marginal distribution, $Lognormal(\mu, \sigma^2)$, as Y by drawing a standard normal z with correlation $\rho_n = ln(1+\rho)$ with $y = \ln(Y)$, and then setting $Z = exp(\mu + \sigma z)$. Correlation ρ_n between the standard normal implies correlation ρ between their exponentiation due to following proof: Algebraic manipulation of the definition of covariance and the lognormal imply $cov(Y,Z) = E[\exp(y + z)] - E[Y]^2$. y and z are both normal, so $y + z$ is normal and $\exp(y + z)$ is lognormal. Using the above covariance formula, the formulas for the mean and standard deviation of a sum correlated normals, the definitions of mean and standard deviation of a lognormal, and the definition of correlation, the implied correlation between y and z is $\rho_n = ln(1+\rho)$.

$$\omega_{ij}^{QUANTILE} = \frac{\left(1 - S_{ij}\left(\delta_{ij}^{QUANTILE}\right)\right) S_{ij}\left(\delta_{ij}^{QUANTILE}\right)\alpha_{QUANTILE}(y_i)}{\int \left(1 - S_{ij}\left(\delta_{ij}^{QUANTILE}\right)\right) S_{ij}\left(\delta_{ij}^{QUANTILE}\right)\alpha_{QUANTILE}(y_i)\partial F(y_i)}. \quad (26)$$

I estimate two versions of this specification. The first uses quintiles because five income bins are common for applications.[58] The second uses deciles, which will show how much estimate improve as the binning becomes more granular.

6.5 Product-Market Coefficients Specification

An alternative to the simple logit that avoids the difficulties in incorporating varying price sensitivity but allows δ_{ijt} to capture taste variation due to price would be a model with product-market specific coefficients on y_i. The associated index is

$$\delta_{ijt}^{PMC} = \beta_j^{PMC} + \alpha^{PMC} P_{jt} + \psi_{jt}\dot{y}_i. \quad (27)$$

A similar demand system was used by the applicants in the FCC's T-Mobile/Sprint merger review process, though that had a more complicated set of covariates.[59]

An advantage of this specification is that it is possible to estimate the choice probabilities perfectly – $\psi_{jt}\dot{y}_i$ captures all the utility variation that is ignored by mis-specifying the relation of price sensitivity to y_i. With strong variation in data on choice and y_i, an econometrician should be able to estimate $\hat{\beta}_j^{PMC} = 5Q_j$, $\hat{\alpha}^{PMC} = -5$, and $\psi_{jt} = -5P_{jt}$: then $\delta_{ijt}^{PMC} = \delta_{ijt}^{TRUE}$. In fact, it would be possible to estimate choice probabilities without prices if there was also a product-market fixed effect.[60] I take this specification to be a "best case" scenario for a mis-specified model that does not allow variation in price sensitivity – there exists an approximation to this implementation in which there is no direct data on y_i but instead data to proxy for the entire value of $\theta y_i P_{jt}$.[61]

The disadvantage of this specification is that the diversion ratios are biased. As the price coefficient is now constant, it factors out in the implied weights:

$$\omega_{ijt}^{CONSTANT} = \frac{\left(1 - S_{ijt}\left(\delta_{ijt}^{PMC}\right)\right) S_{ijt}\left(\delta_{ijt}^{PMC}\right)}{\int \left(1 - S_{ijt}\left(\delta_{ijt}^{PMC}\right)\right) S_{ij}\left(\delta_{ijt}^{PMC}\right)\partial F(y_i)}. \quad (28)$$

Even if the choice probabilities are estimated perfectly, the diversion ratio will still be biased because the weights ignore y_i. In particular, diversion from High Tier products will be overestimated as this formula

[58] Examples with five income groups in a demand system include Goolsbee & Petrin, *supra* note 34; Leemore Dafny & David Dranove, *Do Report Cards Tell Consumers Anything They Don't Already Know? The Case of Medicare HMOs*, 39 RAND J. Econ. 790 (2008); and Amil Petrin & Kenneth E. Train, *A Control Function Approach to Endogeneity in Consumer Choice Models*, 47 J. Marketing Res. 3 (2010).
[59] T-Mobile/Sprint Expert Economic Analysis at 21-24, paras. 48-58.
[60] For example, the authors of the T-Mobile/Sprint demand model did not estimate price sensitivity directly in the demand system but estimated it in an auxiliary procedure. This was because major U.S. mobile telephone services use national pricing, so there is little to no market level price variation. *Id.* at 21, para. 50 & n. 45, and at 63, para. 63 & n. 54.
[61] For example, machine learning techniques could incorporate variation from many variables correlated with $\theta y_i P_{jt}$, or even δ_{ijt} more generally, to produce accurate predictions of S_{ijt}.

gives the price-insensitive consumers who disproportionately choose High Tier products significant weight while in reality they should have little weight at all.

6.6 Random Coefficient Specification

If micro-data on varying taste for price (or other product characteristics) are not available, it is possible to estimate the distribution of y_i by assuming a parametric distribution for it. In this case this amounts to using

$$\delta_{ijt}^{RC} = \beta_j^{RC} + \alpha_i^{RC} P_{jt}. \tag{29}$$

Simulations or numerical integrals of α_i^{RC} are used to form choice probabilities for every guess of the parameters, including those for the distribution of α_i^{RC}. Random coefficient specifications like this can be computationally difficult and time-consuming and are rarely used in merger reviews.

The specification can recreate the true choice probabilities if $\hat{\beta}_j^{RC} = 5Q_j$ and $\hat{\alpha}_i^{RC}$ is estimated to be the negative of a lognormal distribution with $\mu = -0.42 + ln(5)$ and $\sigma = \ 0.92$ (which produces a mean price sensitivity of -5 and skewness of 5). I estimate two versions of this specification: one assuming the correct negative lognormal distribution, and the another assuming a normal distribution. The latter will reflect the impact of mis-specifying the strictly positive and asymmetric price sensitivity distribution with a symmetric distribution over a full support.

To generate the likelihoods for estimation, I use simulation in which a finite number of draws of $\hat{\alpha}_i^{RC}$ are taken and resulting choice probabilities are averaged over the draws.[62] Integration for the market-level diversion ratios and GUPPIs use simulated draws of $\hat{\alpha}_i^{RC}$ using Halton draws, which have good performance in random coefficients models.[63]

6.7 Nested Logit Specification

Nested logit is a popular formulation in discrete choice estimation because it admits more flexible substitution patterns and is relatively simple to estimate.[64] Nested logit is often estimated as a first step in analyzing a demand dataset: if certain estimated parameters are less than one then there is evidence true model lacks IIA.[65] Estimating nested logits is much easier than estimating a random coefficient model because there is no integration of coefficients required. The merger reviews of Aetna/Humana and AT&T/DirecTV used nested logit demand systems for their diversion ratio estimates. Grigolon and Verboeven (2014) compare random coefficient and nested logit estimates in Monte Carlo experiments

[62] I use Stata's asmixlogit function to estimate the random coefficient specifications and use the default settings. This means 50 draws generated through Hammersley sequences. Stata, *asmixlogit — Alternative-Specific Mixed Logit Regression*, https://www.stata.com/manuals15/rasmixlogit.pdf (last visited Nov. 19, 2021).
[63] Kenneth E. Train, *Halton Sequences for Mixed Logit*, 2000 UC Berkeley Working Paper No. E00-278, https://escholarship.org/uc/item/6zs694tp.
[64] *See generally* Train, *supra* note 33, at 77-86.
[65] The parameters in question are the nesting parameters λ_g, which governs substitution patterns as explained later in this subsection. The most well-known formal test is from Jerry Hausman & Daniel McFadden, *Specification Tests for the Multinomial Logit Model*, 52 Econometrica 1219, 1226-1229 (1984).

similar to those in this paper. Using automobile data, they find the nested logit to be a good approximation of random coefficients in substitution patterns.[66]

The nested logit diversion ratio does not share the weighted average formula for individual diversion ratios or weights in (15) and (16). The nested logit varies from the previous logit specifications by assuming that product taste shocks are not i.i.d but are correlated for products in the same product "nest," g, which is one of several mutually exclusive product groupings in the set of nests G. Given the nested logit distribution of taste shocks, the nested logit choice probabilities for a product j in set J^g are

$$S_{ijt} = S_{igt} S_{ij|gt} \tag{30}$$

where

$$S_{igt} = \frac{\left(\sum_{k \in J^g} \exp\left(\frac{\delta_{ikt}}{\lambda_g}\right)\right)^{\lambda_g}}{\sum_{h \in G}\left(\sum_{k \in J^h} \exp\left(\frac{\delta_{ikt}}{\lambda_h}\right)\right)^{\lambda_h}} \tag{31}$$

$$S_{ij|gt} = \frac{\exp\left(\frac{\delta_{ijt}}{\lambda_g}\right)}{\sum_{k \in J^g} \exp\left(\frac{\delta_{ikt}}{\lambda_g}\right)}\,. \tag{32}$$

S_{igt} can be interpreted as the choice probability for a nest and $S_{ij|gt}$ can be interpreted as the choice j probability for a product conditional on the nest g being chosen. Thus this model can be conceptualized as a consumer choosing a nest first according to a logit over the nest-specific "inclusive value," $\sum_{k \in J^g} \exp\left(\frac{\delta_{ikt}}{\lambda_g}\right)$, and then limiting themselves to that nest when making a final decision of the product according to a nest-specific logit.[67] The nest-specific parameters λ_g ("nesting parameters" or "dissimilarity parameters") govern the correlation of the product taste shocks within nests such that as λ_g approaches 0, the product taste shocks become more correlated. As λ_g approaches 1, the shocks become independent and the model reverts to a logit.

Assuming no variation in consumer types, y_i, expected market-level diversion ratios are equal to constant individual diversion ratios:

$$E\left[DR_t^{jk}\right] = DR_{it}^{jk} = -\frac{S_{ikt} S_{ijt}\left(1 + \Lambda_{jkt}\right)}{\frac{S_{ijt}}{\lambda_g}\left(1 - \left(1 - \lambda_g\right) S_{ij|gt} - \lambda_g S_{ijt}\right)}. \tag{33}$$

where Λ_{jk} varies based on whether j and k are in the same nest g:

[66] Laura Grigolon & Frank Verboven, *Nested Logit or Random Coefficients Logit? A Comparison of Alternative Discrete Choice Models of Product Differentiation*, 96 Rev. Econ. Stat. 916 (2014).
[67] It is possible to extend nesting by adding several layers of nests. Train, *supra* note 33, at 86-88.

$$\Lambda_{jkt} = \begin{cases} 0, & j \text{ and } k \text{ are in different nests} \\ \dfrac{\lambda_g^{-1} - 1}{S_{gt}}, & j \text{ and } k \text{ are in the same nest } g \end{cases}.$$

Thus unlike the general model of Section IV where product taste shocks are independent, the individual diversion ratios of the nested logit are not just $\frac{S_{ikt}}{1-S_{ijt}}$. Two products can have high diversion with low shares if λ_g is small – the diversion ratio then approaches $\frac{S_{ik|gt}}{1-S_{ik|gt}}$ as λ_g approaches 1.

My nested logit specification admits no individual level variation, so the δ_{ijt} is the same as the Simple Logit specification:[68]

$$\delta_{ijt}^{NL} = \beta_j^{NL} + \alpha^{NL} P_{jt}. \tag{34}$$

One caveat is that some simulations estimate λ_g to be larger than 1. Borsch-Süpan (1990)[69] shows that this could be consistent with a utility maximizing individual based on other parameters. Kling and Herriges (1995)[70] and Herriges and Kling (1996)[71] find the scope for this possibility is limited in theory and empirical practice, respectively. To sidestep these issues, I constrain λ_g to be at most 1, which I find to be generally within the 95% confidence interval of estimate λ_g when it does violate the bound.[72]

7. Coefficient Estimates

The mean coefficient estimates and the standard deviations for all models are reported in Table 5, along with information on the frequency of positive estimated price sensitivities and McFadden's R^2's.[73] As expected, the Simple Logit without any individual data yields poor estimates – all coefficients are biased towards the origin and the variance explained by the model is low, as reflected in the very low McFadden's R^2. Meanwhile using the correct model with correct data yields highly accurate estimates, with bias of no more than 0.1 and with very small standard deviations. Unsurprisingly, the correct model

[68] When I incorporate individual level data with nested logit find that estimation correctly estimates nesting parameters near 1 when the individual level data explains choices well. Thus, I focus on the case where nesting might best serve as a substitute for when such identifying data is unavailable.
[69] Axel Börsch-Süpan, *On the Compatibility of Nested Logit Models with Utility Maximization*, 43 J. Econometrics 373 (1990).
[70] Catherine L. Kling & Joseph A. Herriges, *An Empirical Investigation of The Consistency of Nested Logit Models with Utility Maximization*, 77 Amer. J. Agric. Econ. 875 (1995).
[71] Joseph A. Herriges & Catherine L. Kling, *Testing the Consistency of Nested Logit Models with Utility Maximization*, 50 Econ. Letters 33 (1996).
[72] I use STATA's *nlogit* command to estimate the nested logit demand system which does not admit inequality constraints. When I find a $\lambda_g > 1$, I re-estimate the nested logit with an equality constraint on the nesting parameter that exceeds 1 and report the corresponding set of estimates. Stata, nlogit — *Nested Logit Regression*, https://www.stata.com/manuals/cmnlogit.pdf (last visited Nov. 19, 2021).
[73] McFadden's $R^2$is a test statistic for likelihood estimates analogous to the R^2: it is 1 minus the ratio of the log likelihood of the model over the log likelihood of a version of the model with only choice-specific intercepts as covariates. It represents how much unexplained variation in the simpler model is explained by the more complicated one. Daniel McFadden, *Conditional Logit Analysis of Qualitative Choice Behavior*, in Frontiers in Econometrics 122 (ed. Paul Zarembka 1973).

has an excellent fit with a McFadden's R^2 of 0.35.[74] However, 34% of samples have some consumers with positive price sensitivities, though in those samples they only make up 2.8% of the sample on average.

The mis-measurement of the micro-data has a very striking effect – estimation performance rapidly declines with modest drops in the correlation between the true and mismeasured data. Even the 0.95 Correlation specification has biases toward the origin of more than 1.0 for every parameter. The 0.75 Correlation specification has biases to toward the origin by a factor of at least 2, and the mean McFadden's R^2 of 0.15 is less than half of the mean for the correct model. Interestingly, the mismeasured data specifications have a smaller prevalence of positive price sensitivities because the attenuation is stronger for the price-centered y_i interactions than for the price coefficient – positive price coefficients happen for only 2.8% of the experiments for 0.95 correlation, and simply do not happen for 0.85 or 0.75 correlation.

The Quintile Coefficients specification demonstrates severe problems with positive price sensitivities. While the fit is excellent with a mean McFadden's R^2 of 0.28, 99% of the experiments exhibit positive price sensitivities for some consumers, with either the lowest and sometimes the second-lowest quintile coefficients being positive. This is because the model is not flexible enough to capture the extreme price sensitivity of consumers with high y_i. Instead, the model compensates by estimating lower quality for every good. However, because High Tier products appear to have worse quality, the only way the specification can rationalize price-insensitive consumers choosing expensive High Tier products is to assign them positive value for price. The implied price sensitivity results in an average price coefficient over all simulations of -2.0 instead of the true -5.0. In contrast, the Decile Coefficients specification is flexible enough so this problem is far less severe. While the quality estimates are still biased downwards, only 34% of simulations have some consumers with positive price sensitivity, and in these simulations these consumers usually only make up the lowest y_i decile. Compared to the Quintile Coefficients specification the average price coefficient over all simulations is still somewhat attenuated but greater at -3.7. The specification has an excellent fit with a mean McFadden's R^2 of 0.33.

The Product-Market Coefficient specification estimates also shows excellent fit. Moreover, the price and quality coefficients are very close to the true values and on average the specification has a slightly higher mean McFadden's R^2 than the correct specification. This is not surprising, as the product-market-specific y_i coefficients are so flexible that they absorb some of the simulation error in choices, i.e. this specification slightly overfits relative to the true specification.

[74] Values of 0.2 to 0.4 represent "excellent fit." Daniel McFadden, *Quantitative Methods for Analyzing Travel Behaviour on Individuals: Some Recent Developments*, in Behavioural Travel Modelling 306 (eds. David Hensher & Peter Stopher 1979).

Table 5: Coefficient Estimates

Parameter	True Values	(1) Pure	(2) Correct	(3) 0.95 Corr. y_i	(4) 0.85 Corr. y_i	(5) 0.75 Corr. y_i	(6) Quantile Coefficients	(7) Decile Coefficients	(8) Product-Market Coefficients	(9) Lognormal Random Coefficient	(10) Normal Random Coefficient	(11) Nested Logit
α	-5	-1.33	-5.00	-3.58	-2.43	-1.98			-5.03			-1.30
		(0.18)	(0.42)	(0.33)	(0.27)	(0.24)			(0.77)			(0.18)
γ	-5		-4.99	-3.29	-1.89	-1.30						
			(0.14)	(0.11)	(0.07)	(0.06)						
μ	1.19									1.20	-1.49	
										(0.24)	(0.40)	
σ	0.92									0.93	1.35	
										(0.14)	(0.61)	
λ_{Low}	1											0.63
												(0.10)
λ_{High}	1											0.58
												(0.11)
Q_1	12.5	2.64	12.50	8.62	5.57	4.38	7.35	10.63	12.61	13.11	4.17	2.81
		(0.25)	(0.60)	(0.48)	(0.39)	(0.35)	(0.48)	(0.55)	(1.09)	(2.94)	(1.27)	(0.25)
Q_2	12.5	2.66	12.50	8.63	5.58	4.40	7.36	10.63	12.61	13.12	4.18	2.81
		(0.25)	(0.60)	(0.49)	(0.39)	(0.35)	(0.49)	(0.55)	(1.09)	(2.93)	(1.27)	(0.25)
Q_3	15	3.67	15.00	10.55	7.01	5.65	8.25	12.66	15.13	15.70	5.15	3.84
		(0.40)	(0.89)	(0.74)	(0.60)	(0.54)	(0.76)	(0.83)	(1.67)	(3.36)	(1.47)	(0.41)
Q_4	15	3.68	15.00	10.56	7.03	5.67	8.26	12.66	15.12	15.70	5.16	3.84
		(0.40)	(0.89)	(0.74)	(0.59)	(0.54)	(0.76)	(0.82)	(1.67)	(3.34)	(1.46)	(0.40)
% Sims w. + Price Sens.		0%	27%	11%	0%	0%	99%	34%	0%	0%	100%	0%
Mean % Sample w. + Price Sens. If Any in Simulation		.	2.8%	2.8%	-	-	26.8%	11.0%	-	-	12.9%	-
McFadden's Pseudo-R^2		0.00	0.35	0.28	0.20	0.15	0.28	0.33	0.35	0.35	0.01	0.01
		(0.00)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.00)	(0.00)

Out of the 100 simulations, means and standard deviations of coefficient estimates and McFadden's Pseudo-R^2 reported. Also reported is the percentage of simulations with some consumer with positive price sensitivity and, conditional on being one of those simulations, the mean % of consumers with positive price sensitivity. Quantile price coefficients of specifications (6) and (7) and the product-market y_i coefficients of specification (8) are too numerous to report here but are available upon request. The μ and σ estimates of specification (10) are not directly comparable to that of (9) but are reported on the same rows to conserve space.

The performance of the Random Coefficient specifications depends on the assumed distribution of the random coefficients. When the correct lognormal distribution is assumed, the corresponding distribution parameters mirrors the true distribution very closely and only slightly overestimates the quality parameters. In contrast, if the normal distribution is assumed, the results are very similar to the Quintile Coefficients specification. The Normal Random Coefficients specification cannot rationalize the long tail of very price-sensitive consumers, and so compensates by assuming worse quality and a large percentage of consumers with positive price sensitivity in *every* simulation: 12.9% on average. The average of normal random price coefficient is also only -1.5, which is far lower than the true mean of -5.0. For both distributional assumptions, the mean McFadden's R^2 is small with 0.007 for Lognormal and 0.005 for Normal. This is a limitation of assuming the coefficients are unobservable: each consumer is observationally identical aside from choice, so a random coefficient choice probability is the average over all draws of the random coefficient, i.e., the same number. Thus every consumer gets the same choice probabilities, and so the likelihood is low because there is never a case where the choice probability of an observed choice is especially high.

The Nested Logit specification has nesting parameters of around 0.6, representing significant nesting. Markets without certain products provide ample evidence that consumer prefer products in the same tier: under the true model consumers without access to one product in a tier disproportionately choose the remaining product in that tier. However, with a constant price coefficient, the specification is unable to infer that this diversion pattern is because of variance in price sensitivity and the difference in tier price. As a result, the non-nesting parameters are biased towards the origin much like the Simple Logit specification. Technically, (33) and (34) imply the effective values for these parameters should be divided through by the appropriate nesting parameter, but even with this correction the coefficients are less than half of the true coefficients. Like the Random Coefficients specification, the choice probabilities are the same across all consumers, so the McFadden's R^2 is small as well at 0.006.

8. Diversion Ratios Results

The bias (mean error), the standard deviation of error, and the root mean square error (RSME) of the expected market-level diversion ratios are reported in Table 6 for all models. In general, precision of the estimates is relatively high with most standard deviations below 0.03. As a result, the RSME corresponds mostly to bias when the bias is non-negligible.

As it returns share-proportional diversion ratios, the Simple Logit specification underestimates diversion within tier. Likewise, the Simple Logit overestimates the Low Tier to High Tier diversion ratios, but actually underestimates the diversion from High Tier to Low Tier. High-to-Low diversion is displaced by the large overestimate of High-to-Outside Option diversion (mean diversion ratio of 0.15 versus true diversion of essentially 0). In contrast, using the Correct Model specification results in biases lower than 0.004 for every different case of diversion.

Table 6: Simulated Diversion Ratios

	Specification	Statistic	From 1 to 0	From 1 to 2	From 1 to 3	From 3 to 0	From 3 to 1	From 3 to 4
	Truth		***0.12***	***0.59***	***0.15***	***0.00***	***0.31***	***0.38***
(1)	*Simple Logit*	*Bias*	0.04	-0.30	0.13	0.15	-0.02	-0.11
		SD	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
		RMSE	[0.05]	[0.30]	[0.13]	[0.15]	[0.03]	[0.11]
(2)	*Correct Model*	*Bias*	0.00	0.00	-0.00	0.00	0.00	-0.00
		SD	(0.02)	(0.02)	(0.01)	(0.00)	(0.02)	(0.03)
		RMSE	[0.02]	[0.02]	[0.01]	[0.00]	[0.02]	[0.03]
(3)	*0.95 γ_i Corr.*	*Bias*	0.03	-0.11	0.04	0.01	0.01	-0.04
		SD	(0.02)	(0.02)	(0.01)	(0.00)	(0.01)	(0.02)
		RMSE	[0.04]	[0.11]	[0.04]	[0.01]	[0.02]	[0.05]
(4)	*0.85 γ_i Corr.*	*Bias*	0.06	-0.21	0.07	0.04	0.01	-0.06
		SD	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
		RMSE	[0.07]	[0.21]	[0.07]	[0.04]	[0.01]	[0.07]
(5)	*0.75 γ_i Corr.*	*Bias*	0.07	-0.25	0.09	0.07	0.00	-0.08
		SD	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
		RMSE	[0.07]	[0.25]	[0.09]	[0.07]	[0.01]	[0.08]
(6)	*Quintile Coefficients*	*Bias*	0.27	-0.14	-0.07	0.03	0.12	-0.29
		SD	(0.04)	(0.02)	(0.02)	(0.30)	(2.07)	(4.38)
		RMSE	[0.27]	[0.14]	[0.07]	[0.30]	[2.06]	[4.37]
(7)	*Decile Coefficients*	*Bias*	0.01	0.00	-0.01	0.00	0.03	-0.06
		SD	(0.02)	(0.02)	(0.01)	(0.00)	(0.03)	(0.05)
		RMSE	[0.02]	[0.02]	[0.01]	[0.00]	[0.04]	[0.08]
(8)	*Product-Market Coefficients*	*Bias*	-0.06	-0.10	0.08	-0.00	-0.06	0.13
		SD	(0.01)	(0.02)	(0.02)	(0.00)	(0.02)	(0.03)
		RMSE	[0.06]	[0.10]	[0.08]	[0.00]	[0.07]	[0.13]
(9)	*Lognormal Random Coefficient*	*Bias*	0.00	0.01	-0.00	0.00	-0.00	0.01
		SD	(0.02)	(0.06)	(0.02)	(0.00)	(0.02)	(0.03)
		RMSE	[0.02]	[0.06]	[0.02]	[0.00]	[0.02]	[0.03]
(10)	*Normal Random Coefficient*	*Bias*	0.13	-0.19	0.03	0.15	0.33	-0.82
		SD	(0.03)	(0.05)	(0.03)	(0.09)	(2.09)	(4.23)
		RMSE	[0.14]	[0.20]	[0.04]	[0.17]	[2.11]	[4.29]
(11)	*Nested Logit*	*Bias*	-0.02	-0.09	0.06	0.09	-0.11	0.13
		SD	(0.01)	(0.06)	(0.02)	(0.01)	(0.03)	(0.07)
		RMSE	[0.02]	[0.11]	[0.06]	[0.09]	[0.11]	[0.15]

Out of 100 simulations, the bias, the standard deviation (SD) of error and the root square error (RMSE) reported.

Unsurprisingly, the measurement error specifications decline in performance as the correlation declines. The 0.95 Correlation Specification moderately overestimates across-tier expected diversion and moderately underestimates within-tier diversion. As the correlation becomes weaker, the specifications yield more inaccurate diversion ratios. In particular, the mean within-Low Tier and within High Tier diversion ratios of the 0.75 correlation specification are about half and three fourths of true diversion ratios, respectively (0.34 vs 0.59; 0.31 vs. 0.38). In contrast, the mean Low-to-High diversion ratio is about 60% higher than under true expected diversion (0.24 vs. 0.15). Looking at Figures 6 and 7, the 0.75 Correlation y_i specification too weakly differentiates consumers in their individual diversion ratios and weights. The distributions of both individual diversion ratios and weights are too flat, and the estimates are quite noisy on an individual basis. The 0.95 and 0.85 Correlated y_i specifications exhibit similar patterns but are less severe.

The performance of the Quantile Coefficients specification depends on the number of quantiles. The downward bias in the quality variables is severe enough that model biases down diversion from the Low Tier (-0.14 within-Low Tier and to -0.7 to High Tier) and biases up diversion to the Outside Option (0.27 to the Outside Option). The bias of within-High Tier diversion ratio is high and negative (-0.29) not because the model underestimates the magnitude of diversion, but because High Tier consumption is predicted to *increase* when price goes up because of the prevalence of positive price sensitivities. However, the precision of this diversion ratio estimate is perhaps even more concerning than the bias. The consumers with positive price sensitivities are more likely to choose the High Tier Products (with high prices) so diverting consumers are mix of price-hating and price-loving consumers. In many cases, positive and negative terms in the denominator for the weights described in (26) more or less cancel out, resulting in very small denominators. Moreover, simulation error causes these small denominators vary between positive and negative across simulations. Weights are thus sometimes negative for consumers with positive price sensitivities in some simulations and often very large in magnitude.[75] Figure 9 shows that weights specific for High Tier products vary widely and can be over 200 times that of the average consumer depending on the simulation.[76] This translates into the standard deviations of the within-High Tier and the High-to-Low expected market-level diversion ratios being more than 4.3 and 2.0, respectively.

In contrast, the Decile Coefficient specification estimates very accurate diversion ratios. Most biases are below 0.03 in magnitude, with the least accurate diversion is being the within-High Tier diversion ratio (bias of -0.06). Examination of individual diversion ratios and weights in Figures 10 and 11, respectively, reveal that ten discrete y_i groups are nearly enough groups to approximate the true distributions without making the positive price sensitivities matter too much. The comparison with poor performance of the Quintile Coefficient specification highlights how important it is to have many groups for discretized data.

[75] This is the same potential issue with positive price coefficients discussed for the Correct specification, though for the Correct specification the number of positive price sensitivities estimated turn out to be negligible. *Supra* Section 6.2.

[76] That is, in one of the simulated markets of 500 consumers, the magnitude of the weight would be more than $200 \times \frac{1}{500} = \frac{2}{5}$.

Figure 6: Individual Diversion Ratios of 0.75 Correlated y_i Specification

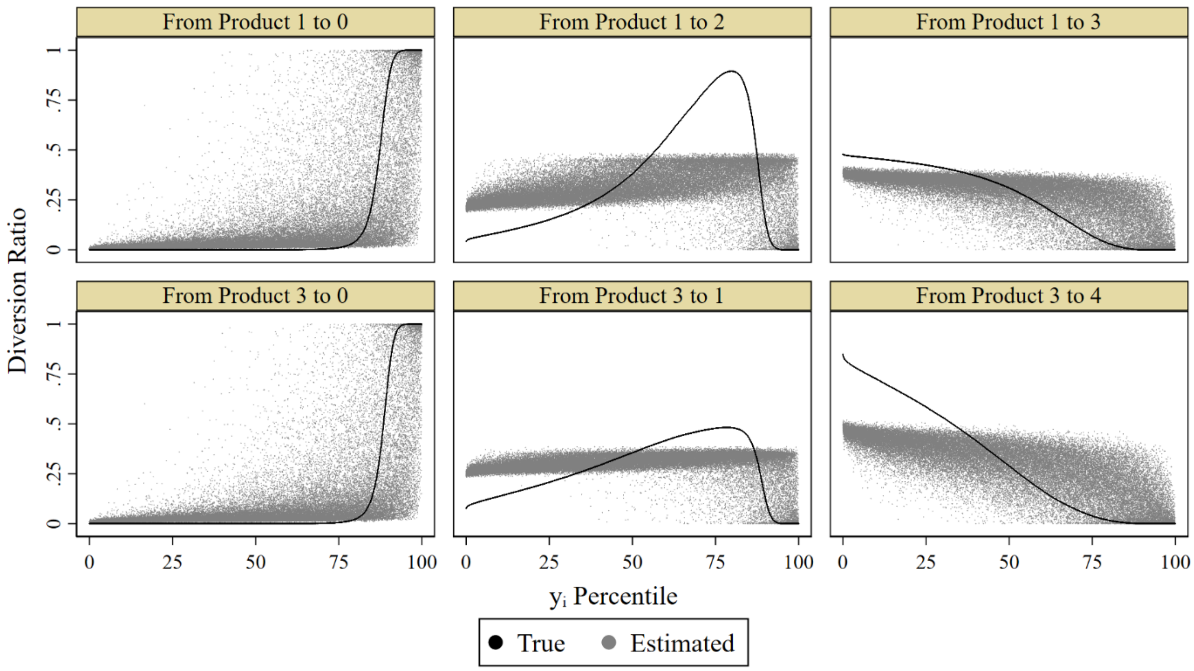

Calculated for all Market 1 individuals across all 100 synthetic datasets.

Figure 7: Weights of 0.75 Correlated y_i Specification

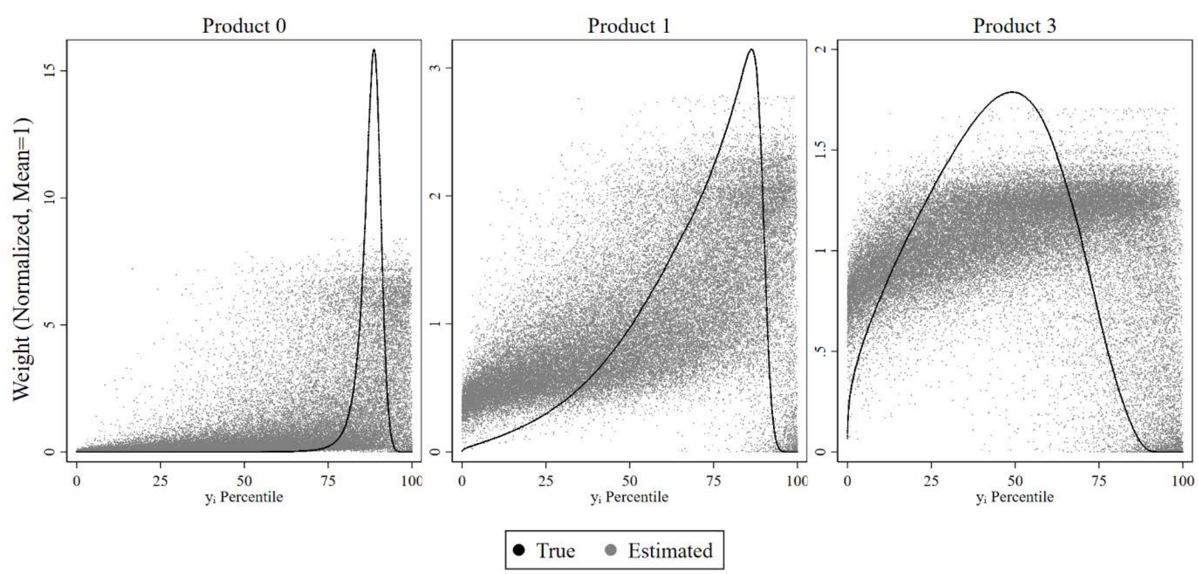

Calculated for all Market 1 individuals across all 100 synthetic datasets. Weights are normalized so that the mean weight is 1 within a single dataset.

Figure 8: Individual Diversion Ratios of Quintile Coefficients Specification

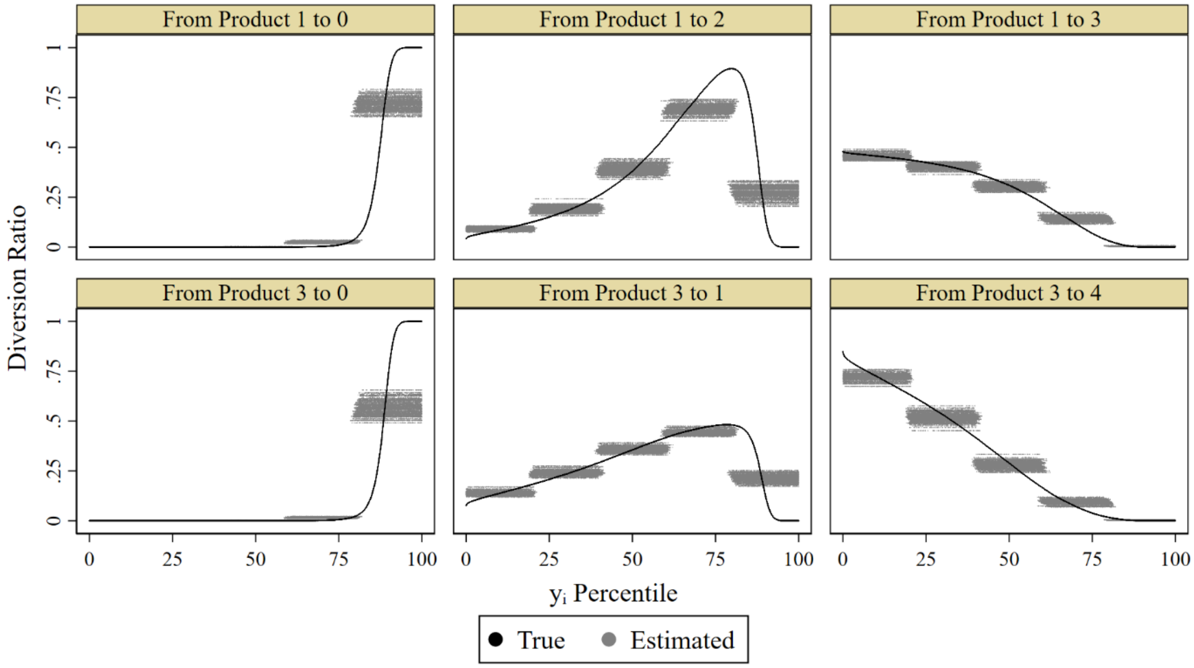

Calculated for all Market 1 individuals across all 100 synthetic datasets.

Figure 9: Weights of Quintile Coefficients Specification

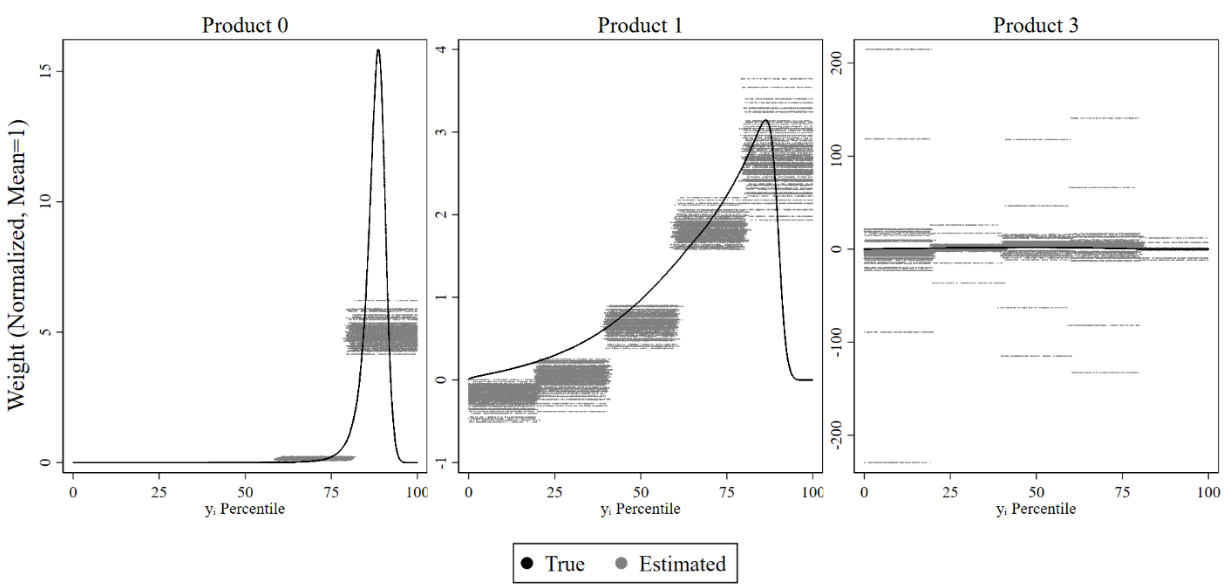

Calculated for all Market 1 individuals across all 100 synthetic datasets. Weights are normalized so that the mean weight is 1 within a single dataset.

Figure 10: Individual Diversion Ratios of Decile Coefficients Specification

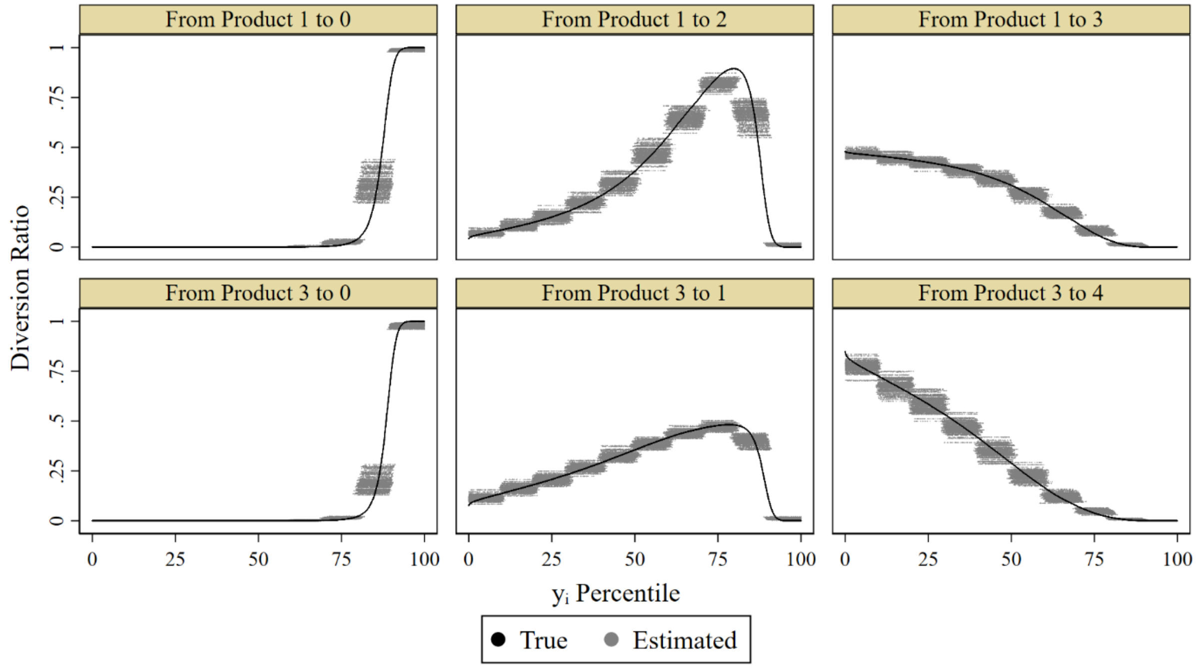

Calculated for all Market 1 individuals across all 100 synthetic datasets.

Figure 11: Weights of Decile Coefficients Specification

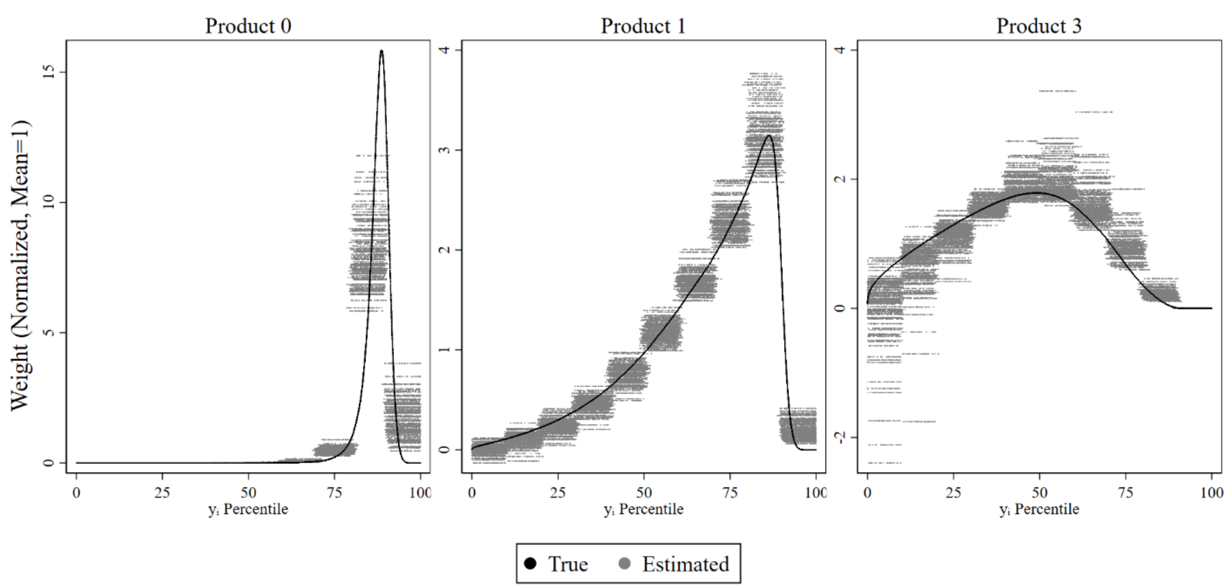

Calculated for all Market 1 individuals across all 100 synthetic datasets. Weights are normalized so that the mean weight is 1 within a single dataset.

The Product-Market Coefficients specification produces biased market-level diversion ratios even though the predicted demand fits "better" than the true model. The mean Low-to-Outside Option diversion ratio is only half as large as they should be (0.06 instead of 0.12). The mean within-Low Tier diversion ratio is about one sixth too low (0.49 instead of 0.59). The mean Low-to-High diversion ratio is about twice as large (0.15 instead of 0.08). The mean High-to-Low diversion ratio is about one sixth too low (0.25 instead of 0.31). Finally, the mean within-High Tier diversion ratio is about one third too large (0.51 instead of 0.38). Examination of the individual diversion ratios in Figure 10 reveals the specification approximates the true individual diversion ratios almost perfectly. The individual diversion ratios in discrete choice models are just functions of choice probabilities, so this is expected from the very good fit of the choice model. The bias that does exist in the market-level diversion ratios comes entirely from the weights, which are mis-specified because they ignore heterogeneity in price sensitivity. Figure 13 shows this specification significantly over-weights diversion from the High Tier, because the specification overpredicts how responsive price-insensitive consumers are to price.

The performance of the Random Coefficient specifications depends on the distribution assumed. As noted before, the Random Coefficient specification cannot generate a high McFadden's R^2 because each estimated individual choice probability is the same. However, the estimated diversion ratios of Lognormal Random Coefficient end up having little bias because it recreates the joint distribution between price sensitivity and choice probabilities needed for the integration over types in (12). The variance in expected market-level diversion ratios is only appreciably larger than using the Correct Model specification for diversion within the Low Tier, where standard deviation is 0.06 instead of 0.02. In contrast, assuming the wrong Normal distribution for the Random Coefficients specification leads to problems similar to the Quintile Coefficients specification, which also estimate a significant amount of positive price coefficients. Attenuated parameters lead to too little diversion within the Low Tier (bias of -0.19) and too much Low-to-Outside option diversion (bias of 0.13). Positive price sensitivity causes even more bias in expected market-level diversion within the High Tier compared to the Quintile Regression (-0.88 vs. -0.29, respectively).[77] Similarly, bias for the High-to-Low diversion ratio is even higher than the Quintile Specification (0.36 vs. 0.12 respectively). These two estimated diversion ratios likewise also have similar imprecision; the within-High Tier diversion ratio has a standard deviation of 4.8 and the High-to-Low diversion ratio has a standard deviation of 2.4.

The Nested Logit specification yields inaccurate diversion ratios, but as demonstrated by the Lognormal Random Coefficient specification this is not due to low McFadden's R^2. The estimated diversion ratios are almost perfectly symmetric – within tier diversion ratios are about 0.5, across tier diversion ratios are about 0.2, and diversion ratios to the Outside Option are about 0.1. As the true diversion ratios vary by tier, this is inaccurate. The biases stem from the specification's estimate of the nesting parameters in both nests to be about equal (about 0.6) and the nest shares also to be about equal (about 45% for the Low Tier and 42% for the High Tier). Because the nest share and nesting parameters are the only things that make the expected market-level diversion ratio vary from the Simple Logit share-proportional diversion ratios, diversion within nests ends up roughly equal as well. In contrast, true diversion is less intense within the High Tier because 1) price sensitive consumers who happen to choose the High Tier are more likely to switch to the Low Tier and 2) High Tier products attract price insensitive

[77] In case of the Normal Random Coefficients, this is caused by very negative outliers, so the median bias of the within-High Tier diversion ratio is actually only -0.22. In the case of the Quintile Regression, the median of the within-High Tier diversion ratio is actually lower than the mean, with a bias of -0.34.

consumer who are less likely react to price changes at all (0.38 within-High Tier vs. 0.59 within-Low Tier diversion ratios).

The Nested Logit results contrast with Grigolon and Verboeven (2014) in which nested logit and random coefficients are good proxies for each other in Monte Carlo experiments.[78] I speculate that this is because the two major categories of true underlying demand that Grigolon and Verboeven (2014) consider are 1) nest-specific random effects, which is similar to nested logit but not to my "characteristics" setup, and 2) there is both a random coefficient and nesting, so including nesting is always important.[79] Thus the poor performance here of Nested Logit should be thought of as emphasizing the importance of misspecification – with a different true demand system the Nested Logit would perform much better.

Figure 12: Individual Diversion Ratios of Product-Market Coefficients Specification

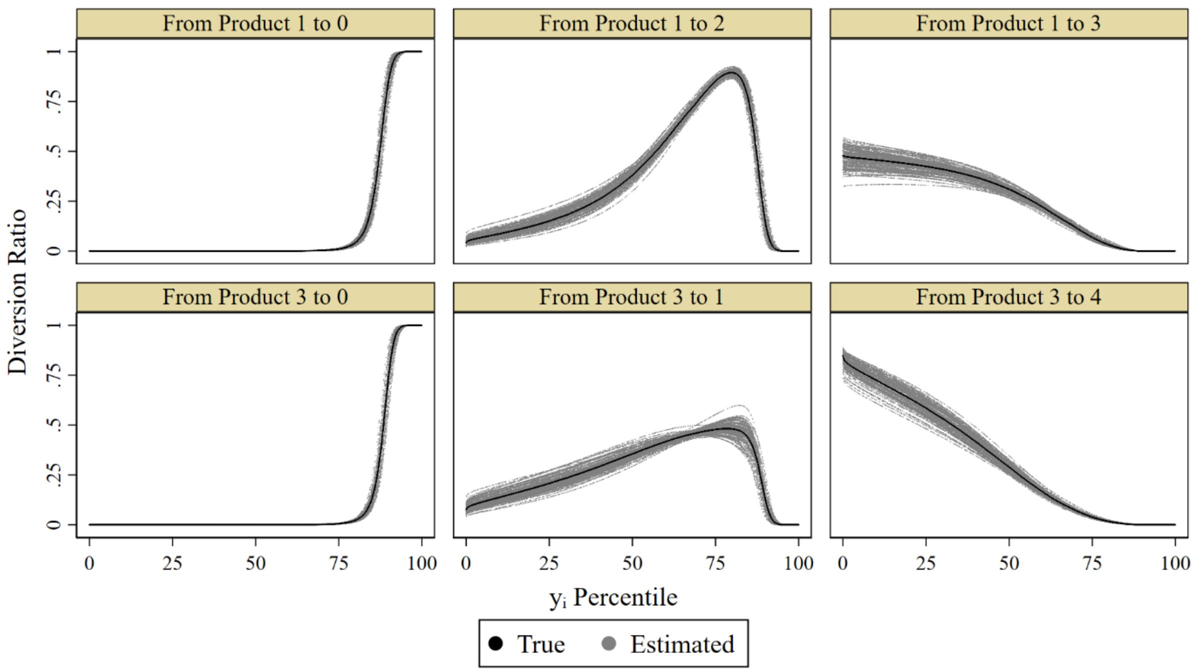

Calculated for all Market 1 individuals across all 100 synthetic datasets.

[78] Grigolon & Verboven, *supra* note 66.
[79] *Id.*

Figure 13: Weights of Product-Market Coefficients Specification

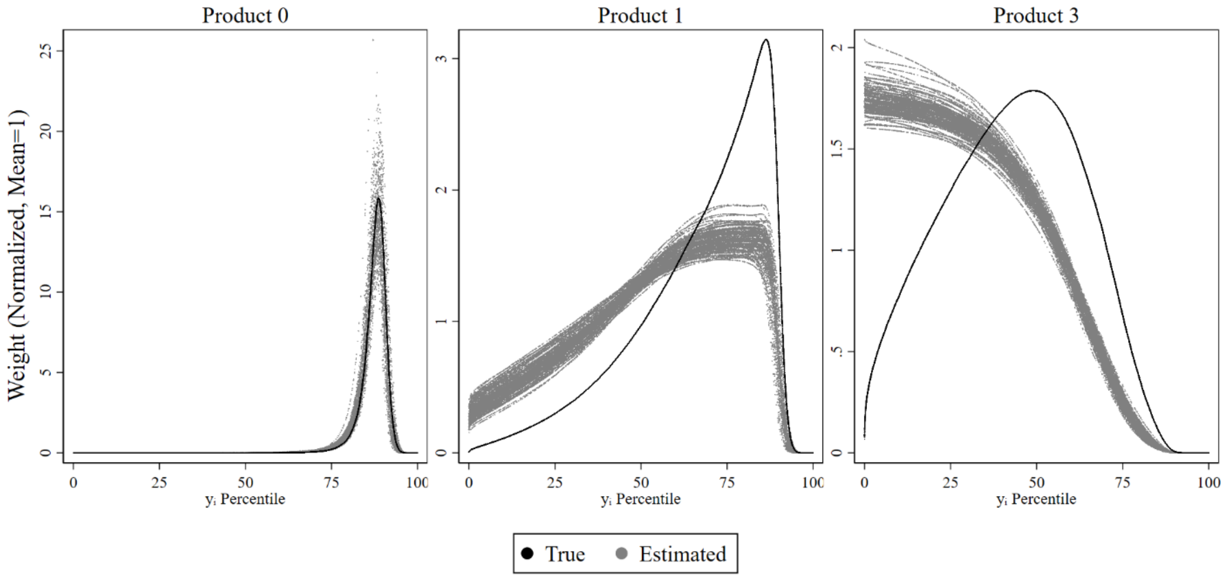

Calculated for all Market 1 individuals across all 100 synthetic datasets. Weights are normalized so that the mean weight is 1 within a single dataset.

Figure 14: Individual Diversion Ratios of Normal Random Coefficient Specification

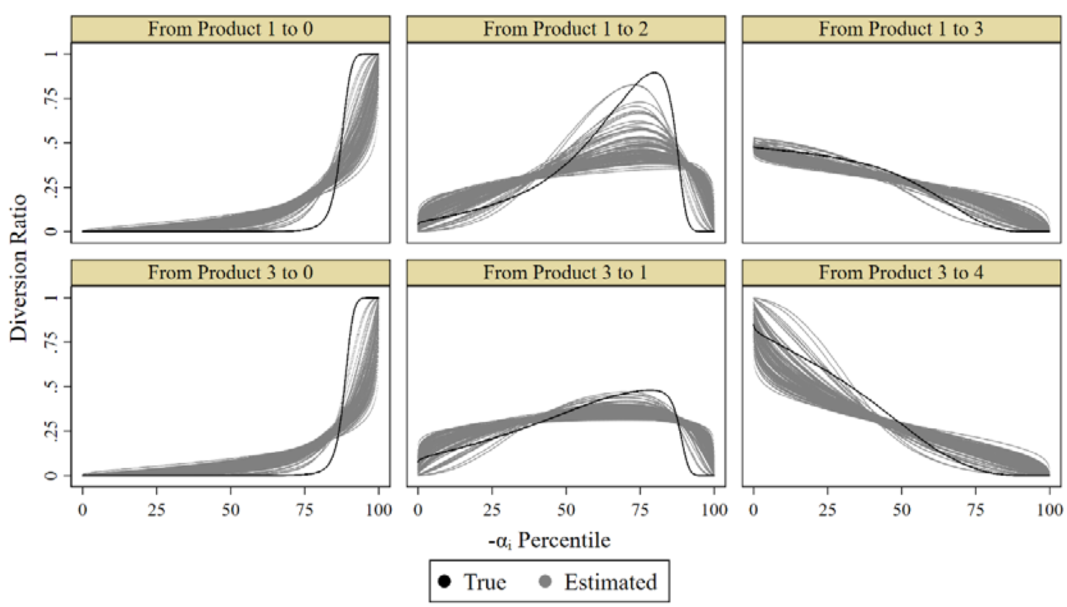

Estimated points calculated for 1,000 draws of $\hat{\alpha}_i$ from the estimated normal distribution for one Market 1 individuals per all 100 synthetic datasets $-\alpha_i$ percentile is specific to each synthetic dataset. True points calculated for all Market 1 individuals across all 100 synthetic datasets. True points plotted over y_i percentiles because $-\alpha_i = \theta y_i$.

Figure 15: Weights of Normal Random Coefficient Specification

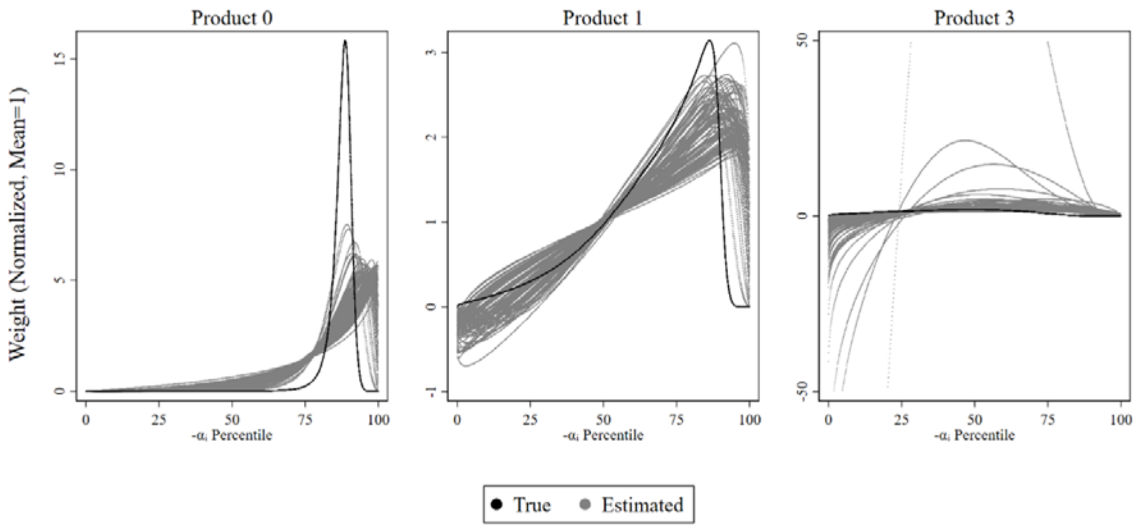

Estimated points calculated for 1,000 draws of $\hat{\alpha}_i$ from the estimated normal distribution for one Market 1 individuals per all 100 synthetic datasets $-\alpha_i$ percentile is specific to each synthetic dataset. True points calculated for all Market 1 individuals across all 100 synthetic datasets. True points plotted over y_i percentiles because $-\alpha_i = \theta y_i$. Weights are normalized so that the mean weight is 1 within a single dataset. Estimated weights for Product 3 censored at $\pm$ 50 which excludes less than 1% of estimated weights.

9. GUPPI Estimates Results

I report the bias, the standard deviation of errors, and the RSME in GUPPIs both calculated using true markups and using markups implied by estimated demand in Table 7. As one might imagine, specifications with accurate coefficients using both true and estimated markups are the ones that produce good estimates of diversion. The Correct specification, the Decile Coefficients specifications, the Lognormal Random Coefficient, and even the 0.95 Correlation y_i specification have low biases. For the Decile Coefficient and 95% Correlation y_i, the bias is greater and positive when using the estimated markups. This is in line with the attenuation of estimated price-sensitivities in these specifications: they underestimate price elasticities and overestimate markups, so GUPPIs are larger.

This is even clearer in specifications that do not have good diversion ratios estimates. When using true markups, the GUPPIs of the Simple Logit mirror the bias of share-based diversion compared to true diversion, with within tier GUPPIs underestimated because they are higher than what market share would suggest. However, the Simple Logit GUPPIs with estimated markups are all overestimated, sometime by more than double. This is because of the very attenuated parameter estimated (-1.3 versus the true -5.0) greatly inflates the markups. By sheer coincidence, the within-High Tier GUPPIs are about right, because the low price parameter counteracts the underestimate of diversion ratios from shares. The 0.85 Correlated y_i and the 0.75 Correlated y_i specifications are similar, in which their GUPPIs biases using observed markups reflect their diversion ratio biases, and the attenuation in price sensitivity leads to large GUPPI upward biases using estimated markups.

Table 7: Simulate GUPPIs

Specification		Statistic	From 1 to 2 True Markup	From 1 to 2 Est. Markup	From 1 to 3 True Markup	From 1 to 3 Est. Markup	From 3 to 1 True Markup	From 3 to 1 Est. Markup	From 3 to 4 True Markup	From 3 to 4 Est. Markup
	Truth		***0.15***		***0.07***		***0.05***		***0.12***	
(1)	*Simple Logit*	*Bias*	-0.07	0.07	0.06	0.13	-0.00	0.09	-0.03	0.01
		SD	(0.01)	(0.03)	(0.01)	(0.04)	(0.00)	(0.02)	(0.01)	(0.02)
		RMSE	[0.08]	[0.08]	[0.06]	[0.13]	[0.00]	[0.09]	[0.03]	[0.02]
(2)	*Correct Model*	*Bias*	0.00	0.00	-0.00	0.00	0.00	0.00	-0.00	0.00
		SD	(0.00)	(0.02)	(0.01)	(0.01)	(0.00)	(0.01)	(0.01)	(0.02)
		RMSE	[0.00]	[0.02]	[0.01]	[0.01]	[0.00]	[0.01]	[0.01]	[0.02]
(3)	*0.95 γ_i Corr.*	*Bias*	-0.03	0.02	0.02	0.03	0.00	0.02	-0.01	0.00
		SD	(0.00)	(0.02)	(0.01)	(0.02)	(0.00)	(0.01)	(0.01)	(0.02)
		RMSE	[0.03]	[0.03]	[0.02]	[0.03]	[0.00]	[0.02]	[0.01]	[0.02]
(4)	*0.85 γ_i Corr.*	*Bias*	-0.05	0.03	0.04	0.07	0.00	0.04	-0.02	0.01
		SD	(0.00)	(0.02)	(0.01)	(0.02)	(0.00)	(0.01)	(0.00)	(0.02)
		RMSE	[0.05]	[0.04]	[0.04]	[0.07]	[0.00]	[0.05]	[0.02]	[0.02]
(5)	*0.75 γ_i Corr.*	*Bias*	-0.06	0.04	0.04	0.09	0.00	0.06	-0.02	0.01
		SD	(0.00)	(0.03)	(0.01)	(0.03)	(0.00)	(0.02)	(0.00)	(0.02)
		RMSE	[0.06]	[0.05]	[0.04]	[0.09]	[0.00]	[0.06]	[0.02]	[0.02]
(6)	*Quintile Coefficient s*	*Bias*	-0.04	0.05	-0.03	0.09	0.02	0.06	-0.09	-19.56
		SD	(0.00)	(0.04)	(0.01)	(1.05)	(0.32)	(0.68)	(1.35)	(199.95)
		RMSE	[0.04]	[0.06]	[0.03]	[1.05]	[0.32]	[0.68]	[1.35]	[199.91]
(7)	*Decile Coefficient s*	*Bias*	0.00	0.03	-0.00	0.02	0.00	0.01	-0.02	0.01
		SD	(0.00)	(0.02)	(0.01)	(0.02)	(0.00)	(0.01)	(0.02)	(0.02)
		RMSE	[0.01]	[0.04]	[0.01]	[0.03]	[0.01]	[0.02]	[0.02]	[0.03]
(8)	*Product-Market Coefficient s*	*Bias*	-0.03	-0.04	0.04	-0.02	-0.01	-0.01	0.04	-0.04
		SD	(0.00)	(0.02)	(0.01)	(0.01)	(0.00)	(0.01)	(0.01)	(0.01)
		RMSE	[0.03]	[0.04]	[0.04]	[0.02]	[0.01]	[0.02]	[0.04]	[0.05]
(9)	*Lognormal Random Coefficient*	*Bias*	0.00	-0.00	-0.00	-0.00	-0.00	-0.00	0.00	0.00
		SD	(0.02)	(0.02)	(0.01)	(0.02)	(0.00)	(0.01)	(0.01)	(0.02)
		RMSE	[0.02]	[0.02]	[0.01]	[0.02]	[0.00]	[0.01]	[0.01]	[0.02]
(10)	*Normal Random Coefficient s*	*Bias*	-0.05	-0.09	0.01	0.22	0.05	0.15	-0.25	-7.50
		SD	(0.01)	(0.05)	(0.02)	(0.55)	(0.33)	(0.37)	(1.30)	(68.34)
		RMSE	[0.05]	[0.11]	[0.02]	[0.55]	[0.33]	[0.39]	[1.32]	[68.41]
(11)	*Nested Logit*	*Bias*	-0.02	0.12	0.03	0.03	-0.02	0.02	0.04	0.04
		SD	(0.02)	(0.04)	(0.01)	(0.03)	(0.00)	(0.02)	(0.02)	(0.02)
		RMSE	[0.03]	[0.13]	[0.03]	[0.04]	[0.02]	[0.03]	[0.05]	[0.05]

Out of 100 simulations, the bias, the standard deviation (SD) of error and the root square error (RMSE) reported.

The Product-Markets Coefficients specification has the reverse bias at work: price insensitive consumers are overly weighted so there is too little diversion to Low Tier products and too much diversion to High Tier products. GUPPIs to Low Tier products are too low and GUPPIs to High Tier products are too high using true markups. However, this bias has a strongly countervailing impact when considering estimated markups – the specification overestimates the price sensitivity of the consumers who do not choose the Outside Option, especially High Tier consumers. Thus, the implied markups are too low, especially for the High Tier. While GUPPIs for diversions to Low Tier goods are only slightly reduced compared to when using the true markups, the GUPPIs for diversions to High Tier goods drop by 50%, so much so that they change from *overestimates* to *underestimates*.

The substantial fraction of price-loving consumers estimated by the Quintile Coefficients and Normal Random Coefficient specifications make their GUPPIs using estimated markups especially volatile. While the GUPPIs using observed markups also are consistent with the biases from the diversion ratios, their attenuated price sensitivities do not always lead to higher GUPPIs using estimated markups. Recall in the case of the Quintile Coefficients specification, price-loving consumers mean that the denominator of the weight formula in (26) is sometime very small and/or negative for the High Tier goods. This value is also the denominator of (21), the formula for estimated markups, so estimated markups can be very large and/or negative. Putting aside the fact negative markups or markups above the High Tier price of 2.17 are nonsensical, these markups in combination with volatile expected market-level diversion ratios leads to highly volatile GUPPIs for diversion to High Tier products. Using estimated markups, the Low Tier to High Tier GUPPI has a standard deviation of bias of 1.35, while the within-High Tier GUPPIs has a standard deviation of almost 200. The Normal Random Coefficient specification has similar issues: the Low Tier to High Tier GUPPI has a standard deviation of 1.3, while the within-High Tier GUPPI has a standard deviation of nearly 70. The precision is somewhat higher than the Quintile Coefficients specification because the Normal Random Coefficient specification estimates fewer positive price sensitive consumers than the Quintile Coefficients specification. While 99% of simulation estimate at least 1 quintile of positive price sensitive consumers, the Normal Random Coefficient specification only estimates 13% on average, so very small denominators of the markup formula (21) are less frequent and not as small.

GUPPIs with true markup for the Nested Logit specification mirror the diversion ratios biases: they underestimate diversion to Low Tier goods and overestimates diversion to High Tier goods. Like the Simple Logit, the price coefficient estimate is attenuated, so markups are too high for Low Tier products but about the right markups for High Tier products. As a result, the Nested Logit GUPPIs using estimated markups are all overestimates of the true GUPPIs.

10. Discussion and Conclusion

This paper documents bias in demand-based diversion ratios that stems from measurement error and misspecification of the demand system. I run Monte Carlo simulations, in which I estimate diversion ratios and GUPPIs with incorrect demand specifications, and compare the results to the true values of the correct demand specification. All of these specifications are commonly used or have been used in practice, and they illustrate different aspects of measurement error and misspecification as sources of bias. Measurement error can bias parameters of the demand system, which will carry over to values calculated using that demand system, including diversion ratios. Misspecification can both lead to biased demand parameters and incorrect formulas for diversion ratios.

I find that even a modest amount of measurement error yields diversion ratios comparable to those obtained from a model accounting for no consumer heterogeneity at all. I find small bias for all diversions ratios in specifications for which the joint distribution of choice probabilities and weights were accurately estimated. The Product-Market Coefficients specification estimates choice probabilities very well, but does not estimate diversion ratios well because it cannot replicate the distribution of weights correctly because it underestimates price sensitivities of price-insensitive consumers. This problem is even worse in the Quintile Coefficients and Normal Random Coefficients specifications, where price insensitive consumers have demand estimates with positive price effects. This results in numerically unstable estimates. GUPPIs replicate these patterns when observed markups are used for calculation, but are biased upwards by the use of estimated markups.

These results reinforce the need for practitioners to verify the robustness of their results through the use of multiple specifications and to ensure data quality. They suggest a few key takeaways for practitioners, especially in merger reviews.

The effectiveness of micro-data on individual characteristics to proxy for differences in price sensitivity can be limited by both measurement error and misspecification. Micro-data used in practice, whether from surveys, imputation or imperfect record-keeping, will often have similar or more extreme amounts of measurement error than the specification of 0.75 correlated data I used. It seems that using consumer data that accurately represent the demographics of its subjects, for example consumer panelist or administrative data, would be completely reliable to proxy for differences in price sensitivity. This also implies a proxies for differences in price sensitivity with weak theoretical bases—for example, using education instead of income because they are correlated—are additionally problematic. Further, even if the data are correct, an incorrect specification like the Product-Market Coefficients specification can produce poor diversion ratio estimates. Such a problem may be hard to notice in practice because the specification can also produce good estimates of demand.

The numerical instability of the Quantile Coefficients specifications suggest that practitioners should use a large number of bins with discretized micro-data to proxying for differences in price sensitivity, but should not use a small number of bins. The Quintile Coefficients specification and Normal Random Coefficient specification also indicate that if one estimates positive price coefficients for some of the consumer sample, the specification may not be flexible enough and should be amended.[80] In general, the instability from positive price sensitivities suggests that demand estimations should ex ante restrict the price sensitivity to be negative.

The difficulties in estimating price sensitivities from data or under parametric assumptions suggest that practitioner should prefer flexible random coefficient models. When the correct distribution of random coefficients is assumed, diversion ratio bias is small and accuracy is high. If practitioners could rely on random coefficient specifications without worrying about misspecification, practitioners could avoid having to deal with the issues involving micro-data entirely. The major caveat is that the misspecification is a large problem with random coefficients: the Normal Random Coefficient specification yielded biased and inaccurate results. Thus practitioners would need specifications that estimate the distribution of the random coefficients or relax the distributional assumptions.

[80] Positive price sensitivities can also be the result of unaddressed endogeneity. Steven T. Berry, *Estimating Discrete-Choice Models of Product Differentiation*, 25 RAND J. Econ. 242, 243, 257-58 (1994).

Several threads in the literature explore these options. The demand estimation procedure of Fox, Kim, Ryan, and Bajari (2014) replaces estimating distribution parameters with estimating the frequency of consumer types.[81] Dubé, Hitsch, and Rossi (2010) estimate a random coefficient model where the random coefficient distribution is a flexible mixture of normal distributions.[82] Brenkers and Verboeven (2006) introduce the Random Coefficient Nested Logit used in the aforementioned Grigolon and Verboeven (2014) which includes random coefficients in a nested logit.[83] Compiani (2022) proposes a nonparametric estimation approach which allows demand beyond standard discrete choice.[84] While these approaches are promising and have prominent applications,[85] they do have more intense data requirements and are more computationally demanding. Such specification would also require addressing common issues I ignored in this paper, including endogeneity and the use of aggregate data. Thus these specifications may be of limited use in a merger review, when time and resources are limited.

Finally, the GUPPI results strongly suggest that practitioners should consider using observed markups for GUPPIs. In my experiments, markups estimated from demand introduce bias that is often more significant than the diversion ratio bias. When trustworthy markup data is available, the practitioner can eliminate this potential channel for measurement error and misspecification bias by using that data instead of estimates from the demand model.

All of the above challenges of properly using demand estimation as the basis of diversion ratio estimates suggest practitioners should carefully consider the use of demand-based diversion ratios over alternatives. As I have shown, measurement error and misspecification introduce potential errors in potentially many different ways for demand-based diversion estimates. Insofar as diversion ratios are indicative of the price predictions from merger simulations based on the same demand system, these results suggest that those price predictions suffer similar biases. When faced with deciding how to estimate diversion ratios, the practitioner should carefully weigh the pros and cons of the demand-based method, especially given the complexity of demand estimation.

While the current study is quite suggestive, it is only done with a specific demand model with a specific set of parameters. The level and direction of biases are specific to the true specification, the kind of the measurement error or misspecification, and the identities of the destination and origin products. I do not claim that the results will fully generalize to all settings. Moreover, I examined a relatively limited set of specifications, and it is unclear what will happen to more complicated specifications. In particular, I did not look at micro-data proxying for sensitivity to product characteristics other than price. In

[81] Jeremy T. Fox, Kyoo Il Kim, Stephen P. Ryan, & Patrick Bajari, *A Simple Estimator for the Distribution of Random Coefficients*, 2 Quantitative Econ. 381 (2011).
[82] Jean-Pierre Dubé, Günter J. Hitsch, & Peter E. Rossi, *State Dependence and Alternative Explanations For Consumer Inertia*, 41 RAND J. Econ. 417 (2010).
[83] Randy Brenkers & Frank Verboven, *Liberalizing a Distribution System: the European Car Market*, 4 J. Eur. Econ. Assoc. 216 (2006); and Grigolon & Verboven, *supra* note 66.
[84] Giovanni Compiani, Market Counterfactuals and the Specification of Multi-Product Demand: A Nonparametric Approach, 13(2) Quantitative Econ. 545 (2022).
[85] Nevo, Turner and Williams (2016) examines residential broadband demand using a frequency-based estimation. Aviv Nevo, John L. Turner, & Jonathan W. Williams, *Usage-Based Pricing and Demand for Residential Broadband*, 84 Econometrica 411 (2016). Random Coefficients Nested Logit is popular in alcoholic beverage demand estimation, where there are widely recognized categories that can serve as nests. *E.g.* Nathan H. Miller & Matthew C. Weinberg, *Understanding the Price Effects of the MillerCoors Joint Venture*, 85 Econometrica 1763 (2017); Eugenio J. Miravete, Katja Seim, & Jeff Thurk, *Market Power and the Laffer Curve*, 86 Econometrica 1651 (2018); and Christopher T. Conlon & Nirupama L. Rao, *Discrete Prices and the Incidence and Efficiency of Excise Taxes*, 12 Amer. Econ. J.: Econ. Pol. 111 (2020).

addition, I assumed extremely favorable conditions for identification, so performance of the low bias specifications may be worse in practical settings. I also do not use real data for my experiments; there may be value in examining the performance of demand models with my data or in experiments with subjects. I therefore look forward to future work that builds on this study.

Works Cited

American Community Survey Office, United States Census Bureau, *AMERICAN COMMUNITY SURVEY 2019 ACS 1-YEAR PUMS FILES ReadMe* 3-4 (2020), https://www2.census.gov/programs-surveys/acs/tech_docs/pums/ACS2019_PUMS_README.pdf.

Simon P. Anderson, Andre De Palma, & Jacques-Francois Thisse, Discrete Choice Theory of Product Differentiation (1992)

Steven T. Berry, *Estimating Discrete-Choice Models of Product Differentiation*, 25 The RAND J. Econ. 242 (1994).

Steven T. Berry & Philip Haile, *Foundations of Demand Estimation*,Handbook of Industrial Organization, Vol. 4, 1 (eds. Kate Ho, Ali Hortaçsu, & Alessandro Lizzeri 2021).

Axel Börsch-Süpan, *On the Compatibility of Nested Logit Models with Utility Maximization*, 43 J. Econometrics 373 (1990).

Randy Brenkers & Frank Verboven, *Liberalizing a Distribution System: the European Car Market*, 4 J. Eur. Econ. Assoc. 216 (2006).

Timothy F. Bresnahan, *Competition and Collusion in the American Automobile Industry: The 1955 Price War*, 35 J. Indus. Econ. 457 (1987).

Yongmin Chen & Marius Schwartz, *Churn Versus Diversion In Antitrust: An Illustrative Model*, 83 Economica 564 (2016).

Giovanni Compiani, *Market Counterfactuals and the Specification of Multi-Product Demand: A Nonparametric Approach*, 13(2) Quantitative Econ. 545 (2022)

Christopher T. Conlon & Julie Holland Mortimer, *Demand Estimation Under Incomplete Product Availability*, 5 Amer. Econ. J.: Microecon. 1 (2013).

Christopher T. Conlon & Julie Holland Mortimer, *Empirical Properties of Diversion Ratios*, 52 RAND J. Econ 693 (2021).

Christopher T. Conlon & Nirupama L. Rao, *Discrete Prices and the Incidence and Efficiency of Excise Taxes*, 12 Amer. Econ. J.: Econ. Pol. 111 (2020).

Philip Crooke, Luke Froeb, Steven Tschantz, & Gregory J. Werden, *Effects of Assumed Demand Form on Simulated Post-Merger Equilibria*, 15(3) Rev. Indus. Org. 205 (1999).

Leemore Dafny & David Dranove, *Do Report Cards Tell Consumers Anything They Don't Already Know? The Case of Medicare HMOs*, 39 RAND J. Econ. 790 (2008).

Jean-Pierre Dubé, Günter J. Hitsch, & Peter E. Rossi, *State Dependence and Alternative Explanations For Consumer Inertia*, 41 RAND J. Econ. 417 (2010).

Kirsten Edwards, *Estimating Diversion Ratios: Some Thoughts on Customer Survey Design, in* European Competition Law Annual 2010: Merger Control in European and Global Perspective 31 (eds. Philip Lowe & Mel Marquis 2013).

Joseph Farrell & Carl Shapiro, *Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition*, 10 BE J. Theoretical Econ. Article number: 00001022021935170415563 (2010).

Joseph Farrell & Carl Shapiro, *Upward Pricing Pressure In Horizontal Merger Analysis: Reply to Epstein and Rubinfeld*, 10 BE J. Theoretical Econ. (2010).

Joseph Farrell & Carl Shapiro, *Upward Pricing Pressure and Critical Loss Analysis: Response*, 2010 CPI Antitrust J.

Jeremy T. Fox, Kyoo Il Kim, Stephen P. Ryan, & Patrick Bajari, *A Simple Estimator for the Distribution of Random Coefficients*, 2 Quantitative Econ. 381 (2011).

Ari D. Gerstle, Helen C. Knudsen, June K. Lee, W. Robert Majure, & Dean V. Williamson, *Economics at the Antitrust Division 2016–2017: Healthcare, Nuclear Waste, and Agriculture*, 51 Rev. Indus. Org. 515 (2017).

Austan Goolsbee & Amil Petrin, *The Consumer Gains from Direct Broadcast Satellites and the Competition With Cable TV*, 72 Econometrica 351 (2004).

Laura Grigolon & Frank Verboven, *Nested Logit or Random Coefficients Logit? A Comparison of Alternative Discrete Choice Models of Product Differentiation*, 96 Rev. Econ. Stat. 916 (2014).

Jerry Hausman, *2010 Merger Guidelines: Empirical Analysis,* 2011 Working Paper, https://economics.mit.edu/files/6603.

Jerry Hausman & Daniel McFadden, *Specification Tests for the Multinomial Logit Model,* 52 Econometrica 1219 (1984).

Joseph A. Herriges & Catherine L. Kling, *Testing the Consistency of Nested Logit Models with Utility Maximization*, 50 Econ. Letters 33 (1996).

Dongling Huang, Christian Rojas, & Frank Bass, *What Happens when Demand is Estimated with a Misspecified Model?* 56 J. Indus. Econ. 809 (2008).

Letter from Maureen Jeffreys, Counsel to AT&T Inc., to Marlene H. Dortch, Secretary, FCC, MB Docket No. 14-90, Attach. (files Jul. 17, 2014), https://ecfsapi.fcc.gov/file/7521680277.pdf.

Catherine L. Kling & Joseph A. Herriges, *An Empirical Investigation of The Consistency of Nested Logit Models with Utility Maximization*, 77 Amer. J. Agric. Econ. 875 (1995).

R. Duncan Luce, Individual Choice Behavior: A Theoretical Analysis 9 (1959).

Daniel McFadden, *Conditional Logit Analysis of Qualitative Choice Behavior,* Frontiers in Econometrics 122 (ed. Paul Zarembka 1973).

Daniel McFadden, *Quantitative Methods for Analyzing Travel Behaviour on Individuals: Some Recent Developments,* in Behavioural Travel Modelling 306 (eds. David Hensher & Peter Stopher 1979).

Nathan H. Miller, Marc Remer, Conor Ryan, & Gloria Sheu, *Upward Pricing Pressure as a Predictor of Merger Price Effects*, 52 Int'l J. Indus. Org. 216 (2017).

Nathan H. Miller & Matthew C. Weinberg, *Understanding the Price Effects of the MillerCoors Joint Venture*, 85 Econometrica 1763 (2017)

Eugenio J. Miravete, Katja Seim, & Jeff Thurk, *Market Power and the Laffer Curve*, 86 Econometrica 1651 (2018).

Serge Moresi & Steven C. Salop, *vGUPPI: Scoring Unilateral Pricing Incentives in Vertical Mergers*, 79 Antitrust L.J. 185 (2013).

Aviv Nevo, *Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry*, 31 RAND J. Econ. 395 (2000).

Aviv Nevo, *Measuring Market Power in the Ready-to-Eat Cereal Industry*, 69 Econometrica 307 (2000).

Aviv Nevo, John L. Turner, & Jonathan W. Williams, *Usage-Based Pricing and Demand for Residential Broadband*, 84 Econometrica 411 (2016).

Daniel P. O'Brien & Steven C. Salop, *Competitive Effects of Partial Ownership: Financial Interest and Corporate Control*, 67 Antitrust L.J. 559 (1999).

Oxera, *Diversion Ratios: Why Does It Matter Where Customers Go if a Shop Is Closed?* Agenda: Advancing Economics in Business (Feb. 15. 2009), https://www.oxera.com/wp-content/uploads/2018/07/Diversion-ratios-updated_1-1.pdf-1.pdf.

Amil Petrin & Kenneth E. Train, *A Control Function Approach to Endogeneity in Consumer Choice Models*, 47 J. Marketing Res. 3 (2010).

Cecilia Rossi, Russell Whitehouse & Alex Moore, *Estimating Diversion Ratios In Hospital Mergers*, 15 J. Competition Law Econ. 639 (2019).

Seth B. Sacher & John Simpson. *Estimating Incremental Margins for Diversion Analysis*, 83 Antitrust L.J. 527 (2020).

Steven C. Salop & Serge Moresi, *Updating the Merger Guidelines: Comments*, 2009 Georgetown Law Faculty Publications and Other Works 1662, https://scholarship.law.georgetown.edu/cgi/viewcontent.cgi?article=2675&context=facpub.

Steven C. Salop, Serge Moresi, & John R. Woodbury, *Scoring Unilateral Effects with the GUPPI: The Approach of the New Horizontal Merger Guidelines*, CRA Competition Memo, Charles River Associates (2010).

Avner Shaked & John Sutton, *Relaxing Price Competition Through Product Differentiation*, 49 Rev. Econ. Stud. 3, (1982).

Carl Shapiro, *Mergers with Differentiated Products*, 10 Antitrust 23 (1996).

Carl Shapiro, *The 2010 Horizontal Merger Guidelines: From Hedgehog to Fox in Forty Years*, 77 Antitrust L.J. 701 (2010).

Carl Shapiro & Howard Shelanski, *Judicial Response to the 2010 Horizontal Merger Guidelines*, 58 Rev. Indus. Org. 51 (2021).

Stata, *asmixlogit — Alternative-Specific Mixed Logit Regression*, https://www.stata.com/manuals15/rasmixlogit.pdf (last visited Nov. 19, 2021).

Stata, *nlogit — Nested Logit Regression*, https://www.stata.com/manuals/cmnlogit.pdf (last visited Nov. 19, 2021).

Kenneth E. Train, *Halton Sequences for Mixed Logit*, 2000 UC Berkeley Working Paper No. E00-278, https://escholarship.org/uc/item/6zs694tp.

Kenneth E. Train, Discrete Choice Methods with Simulation, (2d ed. 2009).

United States Census Bureau, *Explore Census Data*, https://data.census.gov/cedsci/ (last visited Nov. 19, 2021).

Horizontal Merger Guidelines, U.S. Department of Justice and the Federal Trade Commission (April 2, 1992).

Horizontal Merger Guidelines, U.S. Department of Justice and the Federal Trade Commission (Aug. 19, 2010).

Tommaso Valletti & Hans Zenger, *Mergers with Differentiated Products: Where Do We Stand?* 58 Rev. Indus. Org. 179 (2021).

Letter from Nancy Victory, Counsel to T-Mobile, to Marlene H. Dortch, Secretary, FCC, WT Docket No. 18-197, Attach. A (filed Nov. 6, 2018), https://ecfsapi.fcc.gov/file/11060648404338/Nov.%206%20Public%20SuppResponse.pdf.

Gregory J. Werden, *A Robust Test for Consumer Welfare Enhancing Mergers Among Sellers of Differentiated Products*, 44 J. Indus. Econ. 409 (1996).

Robert D. Willig, *Merger Analysis, Industrial Organization Theory, and Merger Guidelines*, 1991 Brookings Papers Econ. Activity: Microeconomics 281.